arXiv:1411.6300v1 [cs.AI] 23 Nov 2014

Discrete Bayesian Networks: The Exact Posterior Marginal Distributions

Do Le (Paul) Minh
Department of ISDS, California State University, Fullerton
CA 92831, USA
dminh@fullerton.edu

November 25, 2014

Abstract: In a Bayesian network, we wish to evaluate the marginal probability of a query variable, which may be conditioned on the observed values of some evidence variables. Here we first present our "border algorithm," which converts a BN into a directed chain. For the polytrees, we then present in details, with some modifications and within the border algorithm framework, the "revised polytree algorithm" by Peot & Shachter (1991). Finally, we present our "parentless polytree method," which, coupled with the border algorithm, converts any Bayesian network into a polytree, rendering the complexity of our inferences independent of the size of network, and linear with the number of its evidence and query variables. All quantities in this paper have probabilistic interpretations.

Keywords: Bayesian networks; Exact inference; Border algorithm; Revised polytree algorithm; Parentless polytree method

1 The Bayesian Networks (BNs)

Consider a directed graph $\mathbb{G}$ defined over a set of ℓ nodes $\mathcal{V} = \{V_1, V_2, ..., V_\ell\}$, in which each node represents a variable. (We denote both a variable and its corresponding node by the same notation, and use the two terms interchangeably.) The pairs of nodes (V_i, V_j) may be connected by either the

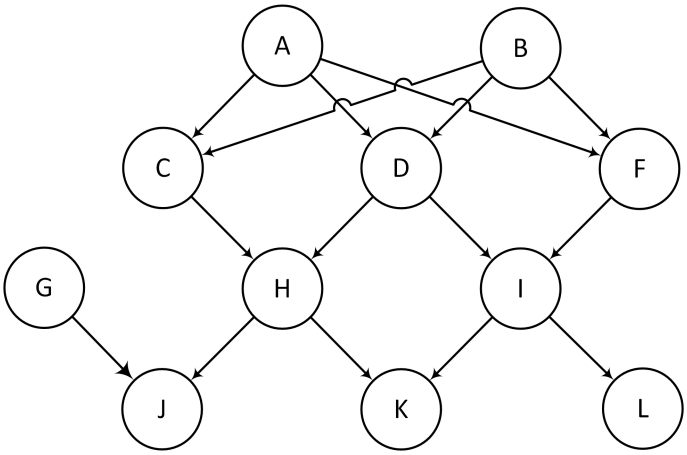

Figure 1: The Bayesian Network $\mathbb{A}$

directed edge $V_i \to V_j$ or $V_j \to V_i$, but not both. It is not necessary that all pairs be connected in this manner. In this paper, we will first use the graph in Figure 1 as an example.

For node $V \in \mathcal{V}$, we call

1. the nodes sending the directed edges to V the "parents" of V. We denote the set of the parents of V by $\mathcal{H}_V$. In Figure 1, $\mathcal{H}_H = \{C, D\}$. A node is said to be a "root" if it has no parents. (For example, nodes A, B and G.)

2. the nodes receiving the directed edges from V the "children" of V. We denote the set of the children of V by $\mathcal{L}_V$. In Figure 1, $\mathcal{L}_D = \{H, I\}$. A node is said to be a "leaf" if it has no children. (For example, nodes J, K and L.) We also call the parents and children of V its "neighbors."

3. the parents of the children of V, except V, the "co-parents" of V. We denote the set of the co-parents of V by $\mathcal{K}_V = \{\cup_{\eta \in \mathcal{L}_V} \mathcal{H}_\eta\} \backslash V$. (We denote by $\mathcal{X} \backslash \mathcal{Y}$ the set $\{X : X \in \mathcal{X}, X \notin \mathcal{Y}\}$. $\mathcal{X} \backslash \mathcal{Y} = \varnothing$ iff $\mathcal{X} \subseteq \mathcal{Y}$.) In our example, $\mathcal{K}_D = \{C, F\}$.

The set of edges connecting nodes V_i and V_j either directly or via other nodes V_k, ..., V_m in the form of $V_i \to V_k \to ... \to V_m \to V_j$ is called a " directed path" from V_i to V_j. We restrict ourselves to the "directed acyclic

graph" (DAG) in which there is no directed path that starts and ends at the same node. If there is a directed path from V_i to V_j, we say V_i is an "ancestor" of V_j and V_j a "descendant" of V_i. Let $\mathcal{N}_V$ and $\mathcal{M}_V$ be the set of all ancestors and descendants of V, respectively. In Figure 1 , $\mathcal{N}_I = \{A, B, D, F\}$, $\mathcal{M}_C = \{H, J, K\}$.

The "Markovian assumption" of a DAG is that every variable is conditionally independent of its non-descendants given its parents. Attached to each node $V \in \mathcal{V}$ is a conditional probability distribution $\Pr\{V|\mathcal{H}_V\}$. If a node has no parent, its distribution is unconditional. We assume in this paper that all $V \in \mathcal{V}$ are discrete, and all conditional probability distributions are in the form of the conditional probability tables (CPTs), taking strictly positive values. We assume that the "size" of $\Pr\{V|\mathcal{H}_V\}$ (that is, the number of possible values of V and $\mathcal{H}_V$) is finite for all $V \in \mathcal{V}$.

A "Bayesian network" (BN) is a pair $(\mathbb{G}, \Theta)$, where $\mathbb{G}$ is a DAG over a set of variables $\mathcal{V} = \{V_1, V_2, ..., V_\ell\}$ (called the "network structure") and Θ a set of all CPTs (called the "network parametrization"). We will refer to the DAG in Figure 1 and its parametrization the Bayesian network $\mathbb{A}$, or the BN $\mathbb{A}$.

It has been shown that the dependence constraints imposed by $\mathbb{G}$ and the numeric constraints imposed by Θ result in the unique joint probability distribution,

$$\Pr\{\mathcal{V}\} = \Pr\{V_1, V_2, ..., V_\ell\} = \prod_{V \in \mathcal{V}} \Pr\{V|\mathcal{H}_V\}. \tag{1}$$

This equation is known as the "chain rule for Bayesian networks" (Pearl, 1987, Equation 3). In our example,

$$\begin{aligned}&\Pr\{A = a, B = b, C = c, ..., L = \ell\}\\ &= \Pr\{A = a\} \Pr\{B = b\} \Pr\{C = c|A = a, B = b\} ... \Pr\{L = \ell|I = i\}.\end{aligned}$$

1.1 The Marginal Distribution

We wish to evaluate the marginal probability $\Pr\{Q\}$, in which $Q \in \mathcal{V}$ is known as a "query variable." This probability may be conditioned on the fact that some other variables in $\mathcal{V}$ are observed to take certain values.

Suppose f is a function defined over a set of variables $\mathcal{X} \subseteq \mathcal{V}$. We say the "scope" of f is $\mathcal{X}$. We list out the scope if necessary, such as $f(\mathcal{X})$; if not, we simply write $f(\cdot)$.

In this paper, suppose $\mathcal{X} = \{\mathcal{Y}, \mathcal{Z}\} \subseteq \mathcal{V}$ where $\mathcal{Y} \cap \mathcal{Z} = \varnothing$ and $\mathcal{Y} = \{Y_1, ..., Y_n\}$. We express $\Pr\{\mathcal{X}\}$ as $\Pr\{\mathcal{Y}, \mathcal{Z}\}$. Given $\Pr\{\mathcal{Y}, \mathcal{Z}\}$, "summing out" (or "eliminating") $\mathcal{Y}$ from $\Pr\{\mathcal{Y}, \mathcal{Z}\}$ means obtaining $\Pr\{\mathcal{Z}\}$ as follows: For every fixed $\mathcal{Z} = z$,

$$\begin{aligned}
&\sum_{\mathcal{Y}} \Pr\{z, \mathcal{Y}\} \\
&= \sum_{Y_1} ... \sum_{Y_{n-1}} \left(\sum_{Y_n} \Pr\{z, Y_1 = y_1, ..., Y_{n-1} = y_{n-1}, Y_n = y_n\} \right) \\
&= \sum_{Y_1} ... \left(\sum_{Y_{n-1}} \Pr\{z, Y_1 = y_1, ..., Y_{n-1} = y_{n-1}\} \right) = \Pr\{z\}.
\end{aligned}$$

We write,

$$\sum_{\mathcal{Y}} \Pr\{\mathcal{Z}, \mathcal{Y}\} = \Pr\{\mathcal{Z}\}. \tag{2}$$

One way to evaluate the marginal probability $\Pr\{V_j\}$ is to use Equation (1) to calculate the joint probability $\Pr\{V_1, ..., V_\ell\}$, then sum out all variables in $\{V_1, ..., V_{j-1}, V_{j+1}, ..., V_\ell\}$. This brute-force method is known to be NP-hard; that is, there is often an exponential relationship between the number of variables ℓ and the complexity of computations (Cooper, 1990). Thus it may be infeasible for large networks.

There have been many attempts in the literature to find the most efficient methods to calculate $\Pr\{Q\}$. They can be divided into two broad categories: the approximate and the exact methods. One example of the approximate methods is using Gibbs samplings to generate "variates" (or "instantiations") for $\mathcal{V}$, then using statistical techniques to find an estimate for $\Pr\{Q\}$. (See Pearl, 1987.) In this paper, we present a method to compute $\Pr\{Q\}$ exactly, apart from precision or rounding errors.

Guo & Hsu (2002) did a survey of the exact algorithms for the BNs, including the two most well-known ones, namely the variable eliminations (Zhang & Poole, 1996; Dechter, 1999) and the clique-tree propagations (Lauritzen & Spiegelhalter, 1988; Lepar & Shenoy, 1999). Other methods reviewed were the message propagations in polytrees (Kim & Pearl, 1983; Pearl 1986a, 1986b), loop cutset conditioning (Pearl, 1986b; Díez, 1996), arc reversal/node reduction (Shachter, 1990), symbolic probabilistic inference (Shachter *et al.*, 1990) and differential approach (Darwiche, 2003). We

also want to mention the more recent LAZY propagation algorithm (Madsen & Jensen, 1999).

In this paper, we first present the *border algorithm*. Like the clique-tree propagation, instead of obtaining the joint probability $\Pr\{V_1, ..., V_\ell\}$, the border algorithm breaks a Bayesian network into smaller parts and calculate the marginal probabilities of these parts, avoiding the exponential blow-ups associated with large networks. In the next section, we first show how a BN can be so divided, in such a way that its independency structure can be exploited. In Section 3, we explain how to calculate the marginal probability of each part when there is no observed evidence. In Section 4, we show how to calculate them, conditional on some observed evidences.

In Section 5, we focus on a special kind of BN called the "polytrees," and present in details, with some modifications and within the border algorithm framework, the "revised polytree algorithm" by Peot & Shachter (1991).

In Section 6, we present our *parentless polytree method*, which, coupled with the border algorithm, can convert any BN into a polytree. This part is static, in that they need to be done only once, off-line, prior to any dialogue with a user. Then we show the dynamic, on-line part of our method, in which the conditional marginal probabilities can be calculated whenever new evidences are entered or queries posed.

Finally, our discussions and summary are presented in Section 7.

2 Partitioning a DAG

In this section, we will show how a BN can be partitioned into smaller parts.

2.1 The Set Relationships

Consider a non-empty set of nodes $\mathcal{X} \subseteq \mathcal{V}$. We also call

1. $\mathcal{H}_\mathcal{X} = \{\cup_{V\in\mathcal{X}}\mathcal{H}_V\}\setminus\mathcal{X}$ the "parent" of $\mathcal{X}$. If $\mathcal{H}_\mathcal{X} = \varnothing$, we say $\mathcal{X}$ is "parentless" (or "ancestral"). For the BN $\mathbb{A}$, $\mathcal{H}_{\{A,H\}} = \{C, D\}$.

2. $\mathcal{L}_\mathcal{X} = \{\cup_{V\in\mathcal{X}}\mathcal{L}_V\} \setminus \{\mathcal{X}, \mathcal{H}_\mathcal{X}\}$ the "child" of $\mathcal{X}$. If $\mathcal{L}_\mathcal{X} = \varnothing$, we say $\mathcal{X}$ is "childless." For the BN $\mathbb{A}$, $\mathcal{L}_{\{A,H\}} = \{J, K\}$. (Although D is a child of A, it is also a parent of H; so it is a member of $\mathcal{H}_{\{A,H\}}$, not of $\mathcal{L}_{\{A,H\}}$.)

3. $\mathcal{K}_\mathcal{X} = \{\cup_{V\in\mathcal{X}}\mathcal{K}_V\}\setminus\{\mathcal{X},\mathcal{H}_\mathcal{X},\mathcal{L}_\mathcal{X}\}$ the "co-parent" of $\mathcal{X}$. If a child of $V\in\mathcal{X}$ is also in $\mathcal{X}$, then all its parents are in $\{\mathcal{X},\mathcal{H}_\mathcal{X}\}$. Thus we are only concerned with the children of V in $\mathcal{L}_\mathcal{X}$. $\mathcal{K}_\mathcal{X}$ therefore can also be defined as $\{\cup_{V\in\mathcal{L}_\mathcal{X}}\mathcal{H}_V\}\setminus\{\mathcal{X},\mathcal{L}_\mathcal{X}\}$. For BN $\mathbb{A}$, $\mathcal{K}_{\{A,H\}} = \{B,G,I\}$. If $\mathcal{K}_\mathcal{X} = \varnothing$, we say $\mathcal{X}$ is "co-parentless."

2.2 The Growing Parentless Set

Consider a parentless set $\mathcal{P}\subseteq\mathcal{V}$. It is "growing" when it "recruits" new members. There are simple algorithms in the literature that allow $\mathcal{P}$ to recruit a member in a "topological order" (that is, after all its parents), so that it is always parentless. (For example, Koller & Friedman, 2009, p. 1146.) Let us call $\mathcal{D} = \mathcal{V}\setminus\mathcal{P}$ the "bottom part" of the BN. We present here an algorithm that not only constructs a growing parentless $\mathcal{P}$, but also divides $\mathcal{P}$ into two parts: $\mathcal{P} = \{\mathcal{A},\mathcal{B}\}$, where $\mathcal{A}$ is called the "top part," and $\mathcal{B}$ the "border" that separates $\mathcal{A}$ from $\mathcal{D}$. It will become clear later why we wish to keep the size of border $\mathcal{B}$ (that is, the number of possible values of $\mathcal{B}$ and $\mathcal{H}_\mathcal{B}$) as small as possible.

We call the members of $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{D}$ the "top variables," the "border variables," and the "bottom variables", respectively.

For the initial top part $\mathcal{A}$, we start with $\mathcal{A} = \varnothing$. We use a co-parentless set of roots as the initial border $\mathcal{B}$. There is at least one such set of roots in a BN. (Suppose a set of roots has a non-root co-parent. Then if we trace through the ancestors of this co-parent, we must encounter another set of roots. Again, if this set is not co-parentless, we trace up further. Eventually, we must see a co-parentless set of roots in a finite BN.) In our example, none of the roots is co-parentless, but the set $\{A,B\}$ is.

All bottom variables will eventually join $\mathcal{P}$. However, we do not choose which bottom variable to join next. This method does not give us control over the membership of the top part $\mathcal{A}$. Instead, we first decide which variable in $\mathcal{B}$ is to be "promoted" to $\mathcal{A}$. The promotion of a variable $B\in\mathcal{B}$ may leave a "hole" in the border $\mathcal{B}$; thus $\mathcal{B}$ no longer separates $\mathcal{A}$ from $\mathcal{D}$. This necessitates recruiting some bottom variables into $\mathcal{B}$ to fill that hole, allowing $\mathcal{P}$ to grow. We call the set of the bottom variables that are recruited into border $\mathcal{B}$ upon the promotion of node B the "cohort" of B and denote it by $\mathcal{C}$. To fill the hole, $\mathcal{C}$ must include at least the part of the "Markov blanket" of B (that is, all its children and co-parents) in $\mathcal{D}$. $\mathcal{C}$ may be empty, or may be more than what we need for $\mathcal{B}$ to separate $\mathcal{A}$ and $\mathcal{D}$.

For $\mathcal{P}$ to remain parentless, it is necessary that $\mathcal{H}_\mathcal{C} \subseteq \mathcal{P}$. Because all members of $\mathcal{C}$ are separated from $\mathcal{A}$ by $\mathcal{B}$, $\mathcal{C}$ cannot have any parent in $\mathcal{A}$. So we only need $\mathcal{H}_\mathcal{C} \subseteq \mathcal{B}$. Below are the many ways by which we can choose the next variable $B \in \mathcal{B}$ to promote to $\mathcal{A}$, approximately in order of preference to keep the size of $\mathcal{B}$ small:

1. B has no bottom children, hence no bottom co-parent. Then $\mathcal{C} = \varnothing$.

2. B has bottom children, but no bottom co-parent. Then $\mathcal{C} = \mathcal{L}_B \cap \mathcal{D}$. (This is why we start $\mathcal{B}$ with a co-parentless set of roots.)

3. B has bottom co-parents, which are roots or have no bottom parents. Then $\mathcal{C} = \{\mathcal{L}_B \cup \mathcal{K}_B\} \cap \mathcal{D}$. In Figure 2, variable J (having co-parent N, with $\mathcal{H}_N = K \notin \mathcal{D}$) can be promoted with cohort $\{N, O\}$. Variable H can also be promoted with cohort $\{M, L\}$, because its co-parent L is a root.

4. B is a fictitious variable $\varnothing$, the cohort of which is a bottom variable having all parents in $\mathcal{P}$. In Figure 2, we can recruit variable V (resulting in new border $\{K, J, I, H, V\}$).

5. B is a fictitious variable $\varnothing$, the cohort of which is a bottom root. In Figure 2, we can recruit root P (resulting in new border $\{K, J, I, H, P\}$) or root S (resulting in new border $\{K, J, I, H, S\}$).

6. B is any variable in $\mathcal{B}$, the cohort of which includes not only its bottom children, but also all their bottom ancestors (hence the bottom co-parents of B and perhaps some roots). In Figure 2, I can be promoted with cohort $\{V, W, U, S\}$.

7. B is a fictitious variable $\varnothing$, the cohort of which includes any bottom variable, together with all its bottom ancestors. Unless it is necessary, the worst (but legal) strategy is to bring all bottom variables simultaneously into $\mathcal{B}$.

8. B is a fictitious variable $\varnothing$, the cohort of which is the whole parentless set $\mathcal{P}^* \subseteq \mathcal{D}$. This is equivalent to "merging" $\mathcal{P}$ and $\mathcal{P}^*$. In Figure 2, we can merge $\mathcal{P} = \{\mathcal{A}, \mathcal{B}\}$ with the parentless $\mathcal{P}^* = \{P, Q, R, S, T, U\}$.

   If $\mathcal{P}^*$ has already been divided into the top part $\mathcal{A}^*$ and the border $\mathcal{B}^*$, then $\mathcal{A}^*$ can be merged with $\mathcal{A}$ and $\mathcal{B}^*$ with $\mathcal{B}$. (This is equivalent to

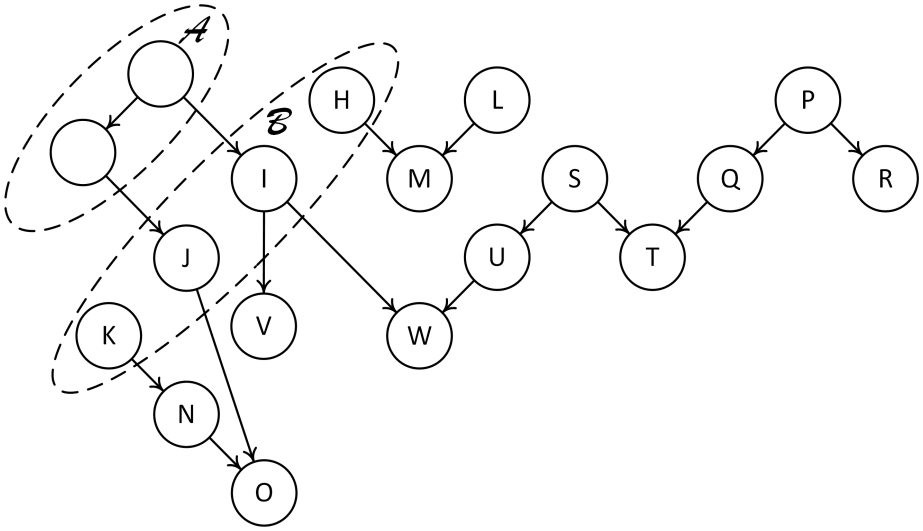

Figure 2: A Bayesian Network

simultaneously promoting all members of $\mathcal{A}^*$ after merging $\mathcal{P}$ and $\mathcal{P}^*$.)
In Figure 2, if $\mathcal{B}^* = \{U, T\}$, then the new border is $\{K, J, I, H, U, T\}$.

We continue to use the notations such as $\mathcal{P}$, $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{D}$. However, we also define the initial top part as $\mathcal{A}_0 = \varnothing$, and denote the initial border, comprising of a co-parentless set of roots, by $\mathcal{B}_0$.

At "time" $j \geq 1$, the variable promoted to $\mathcal{A}_{j-1}$ (which may be a fictitious variable $\varnothing$) is re-named as V_j; the resulting top part becomes $\mathcal{A}_j$. Thus, for all $j \geq 1$,

$$\mathcal{A}_j = \{\mathcal{A}_{j-1}, V_j\} = \{V_1, ..., V_j\}\,. \tag{3}$$

Let $\mathcal{C}_0 = \mathcal{B}_0$. For all $j \geq 1$, the cohort associated with V_j is re-named as $\mathcal{C}_j$. After promoting $V_j \in \mathcal{B}_{j-1}$ and recruiting $\mathcal{C}_j$, the resulting border is, for all $j \geq 1$,

$$\mathcal{B}_j = \{\mathcal{B}_{j-1} \backslash V_j, \mathcal{C}_j\}\,, \tag{4}$$

with $\mathcal{H}_{\mathcal{C}_j} \subseteq \mathcal{B}_{j-1}$.

Let $\mathcal{P}_0 = \mathcal{B}_0$. The parentless set $\mathcal{P}$ grows cohort-by-cohort as, for all $j \geq 1$,

$$\mathcal{P}_j = \{\mathcal{A}_j, \mathcal{B}_j\} = \{\mathcal{A}_{j-1}, V_j\} \cup \{\mathcal{B}_{j-1} \backslash V_j, \mathcal{C}_j\} = \{\mathcal{P}_{j-1}, \mathcal{C}_j\} = \cup_{k=0}^{j} \mathcal{C}_k. \tag{5}$$

Eventually, all variables in $\mathcal{V}$ will join $\mathcal{P}$. Let γ be the time that this happens. We call $\mathcal{B}_\gamma$ the "last" border. Then $\mathcal{V}$ is partitioned into disjoint sets as $\mathcal{V} = \mathcal{P}_\gamma = \cup_{k=0}^{\gamma} \mathcal{C}_k$.

Let the bottom part at time j $(0 \leq j \leq \gamma)$ be

$$\mathcal{D}_j = \mathcal{V} \backslash \mathcal{P}_j = \mathcal{P}_\gamma \backslash \mathcal{P}_j = \cup_{k=j+1}^{\gamma} \mathcal{C}_k = \{\mathcal{C}_{j+1}, \mathcal{D}_{j+1}\}. \tag{6}$$

The above promotion rules do not result in a unique promotion order; and we do not attempt to optimize here, so that the maximum size of all borders $\mathcal{B}_i$ $(i = 1, 2, ..., \gamma)$ is as small as possible. We can heuristically search among all members of $\mathcal{B}_j$ to identify the node whose promotion leads to the smallest next border $\mathcal{B}_{j+1}$, but this does not guarantee a global minimum. The above order of preference may help.

We show the results obtained by one particular promotion order for the BN $\mathbb{A}$ in Table 1. The last column shows the rule we use to promote V_i. The function $\Phi(\mathcal{C}_i)$ will be introduced later.

time i	V_i	$\mathcal{B}_i \backslash \mathcal{C}_i$	$\mathcal{C}_i$	$\Phi(C_i)$	Rule
0	$\varnothing$	$\varnothing$	A, B	$\Pr\{A\} \Pr\{B\}$	
1	A	B	C, D, F	$\Pr\{C\|A, B\} \Pr\{D\|A, B\} \Pr\{F\|A, B\}$	2
2	B	C, D, F	$\varnothing$	1	1
3	C	D, F	H	$\Pr\{H\|C, D\}$	2
4	D	F, H	I	$\Pr\{I\|D, F\}$	2
5	F	H, I	$\varnothing$	1	1
6	H	I	J, K, G	$\Pr\{J\|G, H\} \Pr\{K\|H, I\} \Pr\{G\}$	3
7	G	I, J, K	$\varnothing$	1	1
$\gamma = 8$	I	J, K	L	$\Pr\{L\|I\}$	2

Table 1: The borders obtained for the BN $\mathbb{A}$ from one particular promotion order

To keep $\{\mathcal{P}_i, i = 1, 2, ..., \gamma\}$ parentless, some borders $\mathcal{B}_i$ may have more members than what required to separate $\mathcal{A}_i$ and $\mathcal{D}_i$. Rule 1 is useful in this case, to reduce the membership of $\mathcal{B}_i$ to its minimum. For example, it was used to reduce $\mathcal{B}_4 = \{F, H, I\}$ to $\mathcal{B}_5 = \{H, I\}$.

Here we construct a directed chain of possibly overlapping borders $\{\mathcal{B}_i, i = 1, 2, ..., \gamma\}$, called the "border chain." A border chain is Markovian, in the sense that the knowledge of $\mathcal{B}_j$ is sufficient for the study of $\mathcal{B}_{j+1}$. Figure 3 shows the corresponding border chain for the BN $\mathbb{A}$.

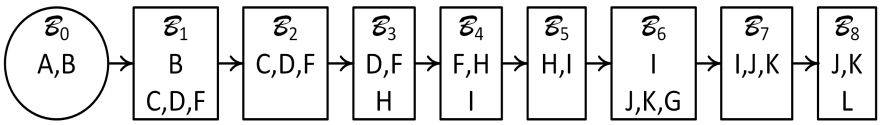

Figure 3: A border chain for the BN $\mathbb{A}$

3 Inferences without Evidences

In this section, we explain how the "prior marginal probability" $\Pr\{\mathcal{B}_i\}$ can be calculated, assuming that no variable is observed taking any value.

3.1 The Parentless Set Probabilities

We first present the following important lemma, which is based on a simple observation that *in a BN, a parentless set of nodes and its parametrization is a Bayesian sub-network*:

Lemma 1 *If $\mathcal{P} \subseteq \mathcal{V}$ is parentless, then*

$$\Pr\{\mathcal{P}\} = \prod_{V \in \mathcal{P}} \Pr\{V|\mathcal{H}_V\}.$$

Proof. The lemma follows from Equation (1) ■

For the BN $\mathbb{A}$, as $\{A, B, D\}$ is parentless,

$$\Pr\{A, B, D\} = \Pr\{A\} \Pr\{B\} \Pr\{D|A, B\}.$$

We do not use $\{A, B, D\}$ however, because D alone does not separate $\{A, B\}$ from the rest of the network. So for a general BN, we start with the parentless $\mathcal{P}_0 = \mathcal{B}_0 = \mathcal{C}_0$ and define

$$\Phi(\mathcal{C}_0) = \Pr\{\mathcal{P}_0\} = \Pr\{\mathcal{B}_0\} = \Pr\{\mathcal{C}_0\} = \prod_{V \in \mathcal{B}_0} \Pr\{V\}. \tag{7}$$

Recall that, when V_j is promoted, it brings a cohort $\mathcal{C}_j$ into $\mathcal{B}_{j-1}$. By the Markovian assumption, we do not need the whole $\Pr\{\mathcal{C}_j|\mathcal{P}_{j-1}\}$, but only $\Pr\left\{\mathcal{C}_j|\mathcal{H}_{\mathcal{C}_j}\right\}$ with $\mathcal{H}_{\mathcal{C}_j} \subseteq \mathcal{B}_{j-1} \subseteq \mathcal{P}_{j-1}$. For all $0 \leq j \leq \gamma$, let us denote the "cohort probability tables" $\Pr\left\{\mathcal{C}_j|\mathcal{H}_{\mathcal{C}_j}\right\}$ by $\Phi(\mathcal{C}_j)$. If $\mathcal{C}_j = \varnothing$, we set $\Phi(\mathcal{C}_j) = 1$. Column 5 of Table 1 shows the cohort probability tables for the BN $\mathbb{A}$.

Theorem 2 *For all* $0 \leq j \leq \gamma$,

$$\Pr\{\mathcal{P}_j\} = \Phi(\mathcal{C}_j) \Pr\{\mathcal{P}_{j-1}\} = \prod_{k=0}^{j} \Phi(\mathcal{C}_k).$$

Proof. From Equation (5),

$$\begin{aligned}\Pr\{\mathcal{P}_j\} &= \Pr\{\mathcal{P}_{j-1}, \mathcal{C}_j\} = \Pr\{\mathcal{C}_j | \mathcal{P}_{j-1}\} \Pr\{\mathcal{P}_{j-1}\} \\ &= \Pr\left\{\mathcal{C}_j | \mathcal{H}_{\mathcal{C}_j}\right\} \Pr\{\mathcal{P}_{j-1}\} = \Phi(\mathcal{C}_j) \Pr\{\mathcal{P}_{j-1}\}.\end{aligned}$$

The theorem follows because $\Pr\{\mathcal{P}_0\} = \Phi(\mathcal{C}_0)$. ■

Theorem 2 can be used to obtain the joint probability of $\mathcal{P}_j$ when j is small. For example,

$$\Pr\{\mathcal{P}_1\} = \Pr\{A, B, C, D, F\} = \Phi(\mathcal{C}_0) \Phi(\mathcal{C}_1).$$

However, as $\mathcal{P}$ grows, eventually we return to Equation (1): $\Pr\{\mathcal{V}\} = \Pr\{\mathcal{P}_\gamma\} = \prod_{k=0}^{\gamma} \Phi(\mathcal{C}_k)$, which is what we did not want to use in the first place.

Fortunately, as we will see in the next section, what we have here is a very "cruel" parentless set of nodes that, after promoting and extracting information from a member, immediately "eliminates" that member!

3.2 The Border Probabilities

We now show how $\Pr\{\mathcal{B}_j\}$ can be recursively calculated from $\Pr\{\mathcal{B}_{j-1}\}$:

Theorem 3 *For all* $1 \leq j \leq \gamma$,

$$\Pr\{\mathcal{B}_j\} = \sum_{V_j} \Phi(\mathcal{C}_j) \Pr\{\mathcal{B}_{j-1}\}.$$

Proof. For all $\mathcal{P}_j$ $(1 \leq j \leq \gamma)$, our strategy is not to eliminate all members of $\mathcal{A}_j$ at the same time, in the form of

$$\Pr\{\mathcal{B}_j\} = \sum_{\mathcal{A}_j} \Pr\{\mathcal{A}_j, \mathcal{B}_j\} = \sum_{\mathcal{A}_j} \Pr\{\mathcal{P}_j\} = \sum_{\mathcal{A}_j} \Phi(\mathcal{C}_j) \Pr\{\mathcal{P}_{j-1}\}.$$

Rather, we eliminate the variables in $\mathcal{A}_j$ one-by-one: After variable V_j is promoted into $\mathcal{A}_j = \{\mathcal{A}_{j-1}, V_j\}$, it is immediately eliminated. In other words,

because the scope of $\Phi(\mathcal{C}_j)$ (which is $\{\mathcal{C}_j, \mathcal{H}_{\mathcal{C}_j}\} \subseteq \{\mathcal{C}_j, \mathcal{B}_{j-1}\}$) does not include any member of $\mathcal{A}_{j-1}$, as far as the summing out of $\mathcal{A}_{j-1}$ is concerned, $\Phi(\mathcal{C}_j)$ can be treated as a constant:

$$\Pr\{\mathcal{B}_j\} = \sum_{\{\mathcal{A}_{j-1}, V_j\}} \Phi(\mathcal{C}_j) \Pr\{\mathcal{P}_{j-1}\} = \sum_{V_j} \Phi(\mathcal{C}_j) \left(\sum_{\mathcal{A}_{j-1}} \Pr\{\mathcal{A}_{j-1}, \mathcal{B}_{j-1}\} \right),$$

hence the theorem. ■

There must be one value of τ $(0 \leq \tau \leq \gamma)$ such that $\Pr\{\mathcal{B}_\tau\}$ can be calculated. At least, from Equation (7), we know $\Pr\{\mathcal{B}_0\}$. Starting with $\Pr\{\mathcal{B}_\tau\}$, we can calculate $\Pr\{\mathcal{B}_j\}$ for all $\tau < j \leq \gamma$ recursively by the above theorem.

We call our algorithm the "border algorithm" because it breaks the large joint probability $\Pr\{\mathcal{V}\}$ down into many smaller border probabilities $\Pr\{\mathcal{B}_j\}$, thus avoiding the exponential blow-ups associated with large networks. That is why we want the size of the largest border to be as small as possible.

We now show how the marginal probabilities can be obtained given some evidences.

4 Inferences with Evidences

For a variable $V \in \mathcal{V}$, let $\mathrm{Va}(V)$ be the set of possible values of V such that $\Pr\{V = v | \mathcal{H}_V\} > 0$. A variable E is said to be an "evidence variable" if it is observed taking value only in a subset $\mathrm{Va}^e(E) \subset \mathrm{Va}(E)$. Variable V is non-evidential if $\mathrm{Va}^e(V) = \mathrm{Va}(V)$. For example, suppose $\mathrm{Va}(X) = \{1, 2, 3\}$. If X is observed not taking value 3, then it is evidential with $\mathrm{Va}^e(X) = \{1, 2\}$. Let $\mathcal{E}$ be the set of all evidence variables.

Consider set $\mathcal{Y} = \{Y_1, ..., Y_n\} \subseteq \mathcal{V}$. We denote the event that $\mathcal{Y}$ occurs by

$$[\mathcal{Y}] = \{\mathcal{Y} \in \mathrm{Va}^e(Y_1) \times ... \times \mathrm{Va}^e(Y_n)\}.$$

If $\mathcal{Y} \cap \mathcal{E} = \varnothing$, $[\mathcal{Y}]$ is a sure event. Thus $[\mathcal{Y}] = [\mathcal{Y} \cap \mathcal{E}]$.

One of the most important tasks in analyzing a BN is to calculate $\Pr\{Q | [\mathcal{E}]\}$, which is known as the "posterior marginal distribution" (or the "conditional marginal distribution") of a "query variable" Q.

For the rest of this paper, we will show how we can first calculate the joint distribution table $\Pr\{Q, [\mathcal{E}]\}$ for all possible values of Q. This allows

us to calculate $\Pr\{[\mathcal{E}]\} = \sum_Q \Pr\{Q, [\mathcal{E}]\}$ and then

$$\Pr\{Q|\,[\mathcal{E}]\} = \frac{\Pr\{Q, [\mathcal{E}]\}}{\sum_Q \Pr\{Q, [\mathcal{E}]\}} = \frac{\Pr\{Q, [\mathcal{E}]\}}{\Pr\{[\mathcal{E}]\}}.$$

4.1 The Evidence Indicator Columns

Consider a table having t rows, each row corresponding to an instantiation of a set of variables $\mathcal{X} \subseteq \mathcal{V}$. If an evidence variable E is in $\mathcal{X}$, we define an "evidence indicator column" $\boldsymbol{I}_E$ having size t, such that it takes value 1 if $E \in \mathrm{Va}^e(E)$, and 0 otherwise. For a non-evidence variable $V \in \mathcal{X}$, we also define the column $\boldsymbol{I}_V$ having size t, all members of which are 1.

We will use the following notation for a set of nodes $\mathcal{X} = \mathcal{Y} \cup \mathcal{Z} \subseteq \mathcal{V}$:

$$\boldsymbol{I}_{\mathcal{X}} = \prod_{V \in \mathcal{X}} \boldsymbol{I}_V = \boldsymbol{I}_{\mathcal{Y} \cup \mathcal{Z}} = \boldsymbol{I}_{\mathcal{Y}} \boldsymbol{I}_{\mathcal{Z}}.$$

Multiplying a table having scope $\mathcal{X}$ with column $\boldsymbol{I}_{\mathcal{X}}$ is equivalent to zeroing out the rows inconsistent with the evidences in $\mathcal{X}$.

For the CPT $\Pr\{V|\mathcal{H}_V\}$, we define its "reduced CPT" as:

$$\Pr_r\{V|\mathcal{H}_V\} = \Pr\{V|\mathcal{H}_V\}\, \boldsymbol{I}_{V \cup \mathcal{H}_V}. \tag{8}$$

Previously we defined the cohort probability tables $\Phi(\mathcal{C}_j) = \Pr\left\{\mathcal{C}_j|\mathcal{H}_{\mathcal{C}_j}\right\}$ for all $0 \leq j \leq \gamma$. We now define the "reduced cohort probability tables" as

$$\phi(\mathcal{C}_j) = \Phi(\mathcal{C}_j)\, \boldsymbol{I}_{\mathcal{C}_j \cup \mathcal{H}_{\mathcal{C}_j}}.$$

The following lemma is the evidential version of Equation (2):

Lemma 4 *Given $\mathcal{Y}, \mathcal{Z} \subseteq \mathcal{V}$, $\mathcal{Y} \cap \mathcal{Z} = \varnothing$, then with scope $\mathcal{Z}$,*

$$\Pr\{\mathcal{Z}, [\mathcal{Y}]\} = \sum_{\mathcal{Y}} \Pr\{\mathcal{Z}, \mathcal{Y}\}\, \boldsymbol{I}_{\mathcal{Y}}.$$

Proof. If $\mathcal{Y} \cap \mathcal{E} = \varnothing$, we have Equation (2) because $\boldsymbol{I}_{\mathcal{Y}} = 1$ and $\Pr\{\mathcal{Z}, [\mathcal{Y}]\} = \Pr\{\mathcal{Z}\}$. Suppose $\mathcal{Y} = \{Y_1, ..., Y_n\}$ is observed taking value in $\mathrm{Va}^e(Y_1) \times ... \times$

$\mathrm{Va}^e(Y_n)$. For every fixed $\mathcal{Z} = z$, summing out $\mathcal{Y}$ yields:

$$\begin{aligned}
&\sum_{\mathcal{Y}} \Pr\{z, \mathcal{Y}\} \boldsymbol{I}_{\mathcal{Y}} \\
&= \sum_{Y_1} ... \sum_{Y_{n-1}} \left(\sum_{Y_n} \Pr\{z, Y_1 = y_1, ..., Y_{n-1} = y_{n-1}, Y_n = y_n\} \prod_{i=1}^{n} \boldsymbol{I}_{Y_i} \right) \\
&= \sum_{Y_1} ... \left(\sum_{Y_{n-1}} \Pr\{z, Y_1 = y_1, ..., Y_{n-1} = y_{n-1}, Y_n \in \mathrm{Va}^e(Y_n)\} \prod_{i=1}^{n-1} \boldsymbol{I}_{Y_i} \right) \\
&= \sum_{Y_1} \Pr\{z, Y_1 = y_1, Y_2 \in \mathrm{Va}^e(Y_2), ..., Y_n \in \mathrm{Va}^e(Y_n)\} \boldsymbol{I}_{Y_1} = \Pr\{z, [\mathcal{Y}]\}.
\end{aligned}$$

■

We are now ready to obtain the necessary information for the calculations of $\Pr\{Q, [\mathcal{E}]\}$.

4.2 The Downward Pass for the Top Evidences

We first consider the "top evidences" within the top part $\mathcal{A}_j$, and define the following notation: For all $0 \leq j \leq \gamma$, by Lemma 4,

$$\begin{aligned}
\Pi(\mathcal{B}_j) &= \Pr\{\mathcal{B}_j, [\mathcal{A}_j]\} \boldsymbol{I}_{\mathcal{B}_j} \\
&= \sum_{\mathcal{A}_j} \Pr\{\mathcal{B}_j, \mathcal{A}_j\} \boldsymbol{I}_{\mathcal{A}_j \cup \mathcal{B}_j} = \sum_{\mathcal{A}_j} \Pr\{\mathcal{P}_j\} \boldsymbol{I}_{\mathcal{P}_j}.
\end{aligned} \tag{9}$$

The following theorem is the evidential version of Theorem 3:

Theorem 5 *For all* $1 \leq j \leq \gamma$,

$$\Pi(\mathcal{B}_j) = \sum_{V_j} \phi(\mathcal{C}_j) \Pi(\mathcal{B}_{j-1}).$$

Proof. Because $\mathcal{H}_{\mathcal{C}_j} \subseteq \mathcal{P}_{j-1}$, we have $\mathcal{P}_j = C_j \cup \mathcal{P}_{j-1} = C_j \cup \mathcal{H}_{\mathcal{C}_j} \cup \mathcal{P}_{j-1}$. From Definition (9) and Theorem 2,

$$\begin{aligned}
\Pi(\mathcal{B}_j) &= \sum_{\mathcal{A}_j} \Pr\{\mathcal{P}_j\} \boldsymbol{I}_{\mathcal{P}_j} = \sum_{\mathcal{A}_j} \Phi(\mathcal{C}_j) \boldsymbol{I}_{\mathcal{C}_j \cup \mathcal{H}_{\mathcal{C}_j}} \Pr\{\mathcal{P}_{j-1}\} \boldsymbol{I}_{\mathcal{P}_{j-1}} \\
&= \sum_{\mathcal{A}_j} \phi(\mathcal{C}_j) \Pr\{\mathcal{P}_{j-1}\} \boldsymbol{I}_{\mathcal{P}_{j-1}}.
\end{aligned}$$

From Equation (3), and as the scope of $\phi(\mathcal{C}_j)$ (which is $\{\mathcal{C}_j, \mathcal{H}_{\mathcal{C}_j}\}$) is not in $\mathcal{A}_{j-1}$,

$$\Pi(\mathcal{B}_j) = \sum_{V_j} \phi(\mathcal{C}_j) \sum_{\mathcal{A}_{j-1}} \Pr\{\mathcal{P}_{j-1}\} \boldsymbol{I}_{\mathcal{P}_{j-1}} = \sum_{V_j} \phi(\mathcal{C}_j) \Pi(\mathcal{B}_{j-1}).$$

■

There must be one value of ν ($0 \leq \nu \leq \gamma$) such that $\Pi(\mathcal{B}_\nu)$ can be calculated. Let α be the first time an evidence variable is recruited into $\mathcal{P}$. For all $0 \leq j < \alpha$, $\mathcal{A}_j$ has no evidence and $\boldsymbol{I}_{\mathcal{B}_j} = 1$; thus $\Pi(\mathcal{B}_j) = \Pr\{\mathcal{B}_j\}$. $\mathcal{A}_\alpha$ also has no evidence and thus,

$$\Pi(\mathcal{B}_\alpha) = \Pr\{\mathcal{B}_\alpha\} \boldsymbol{I}_{\mathcal{B}_\alpha}. \tag{10}$$

Starting with $\Pi(\mathcal{B}_\nu)$, we can calculate $\Pi(\mathcal{B}_j)$ recursively for all $\nu < j \leq \gamma$ by the above theorem.

For the BN $\mathbb{A}$, assume $\mathcal{E} = \{H = h, K = k\}$. Since $H \in \mathcal{C}_3$, $\alpha = 3$. Thus

$$\begin{aligned}
&\Pi(\mathcal{B}_0) = \Pr\{\mathcal{B}_0\} = \Pr\{A, B\};\\
&\Pi(\mathcal{B}_1) = \Pr\{\mathcal{B}_1\} = \Pr\{B, C, D, F\};\\
&\Pi(\mathcal{B}_2) = \Pr\{\mathcal{B}_2\} = \Pr\{C, D, F\};\\
&\Pi(\mathcal{B}_3) = \Pr\{D, F, H\} \boldsymbol{I}_H = \Pr\{D, F, h\}.
\end{aligned}$$

1. Border $\mathcal{B}_4 = \{F, H, I\}$ has $V_4 = D$ and $\mathcal{C}_4 = I$:

$$\begin{aligned}
\Pi(\mathcal{B}_4) &= \sum_D \phi(\mathcal{C}_4) \Pi(\mathcal{B}_3) = \sum_D \Pr\{I|D, F\} \Pr\{D, F, h\}\\
&= \sum_D \Pr\{D, F, I, h\} = \Pr\{F, I, h\}.
\end{aligned}$$

2. Border $\mathcal{B}_5 = \{H, I\}$ has $V_5 = F$ and $\mathcal{C}_5 = \varnothing$:

$$\Pi(\mathcal{B}_5) = \sum_F \phi(\mathcal{C}_5) \Pi(\mathcal{B}_4) = \sum_F \Pr\{F, I, h\} = \Pr\{I, h\}.$$

3. Border $\mathcal{B}_6 = \{I, J, K, G\}$ has $V_6 = H = h$ and $\mathcal{C}_6 = \{J, K, G\}$:

$$\begin{aligned}
\Pi(\mathcal{B}_6) &= \phi(\mathcal{C}_6) \Pi(\mathcal{B}_5)\\
&= \Pr\{J|G, h\} \Pr\{k|I, h\} \Pr\{G\} \Pr\{I, h\} = \Pr\{I, J, G, k, h\}.
\end{aligned}$$

4. Border $\mathcal{B}_7 = \{I, J, K\}$ has $V_7 = G$ and $\mathcal{C}_7 = \varnothing$:

$$\Pi(\mathcal{B}_7) = \sum_G \phi(\mathcal{C}_7)\,\Pi(\mathcal{B}_6) = \sum_G \Pr\{I, J, G, k, h\} = \Pr\{I, J, k, h\}.$$

5. Border $\mathcal{B}_8 = \{J, K, L\}$ has $V_8 = I$ and $\mathcal{C}_8 = L$:

$$\Pi(\mathcal{B}_8) = \sum_I \phi(\mathcal{C}_8)\,\Pi(\mathcal{B}_7) = \sum_I \Pr\{L|I\}\Pr\{I, J, k, h\} = \Pr\{J, L, k, h\}.$$

4.3 The Upward Pass for the Bottom Evidences

Moving downward border-by-border from $\mathcal{B}_0$ to $\mathcal{B}_j$, we can only collect information about the top evidences inside $\mathcal{P}_j$. To collect information about the bottom evidences inside $\mathcal{D}_j$, we move upward from the last border $\mathcal{B}_\gamma$ to $\mathcal{B}_j$.

In the downward passes, we make use of the parentless property of $\mathcal{P}_j$; in the upward passes we need the fact that the border $\mathcal{B}_j$ separates $\mathcal{A}_j$ and $\mathcal{D}_j = \mathcal{V}\backslash\mathcal{P}_j$. Thus, to study $\mathcal{D}_j$, we do not need the information of the whole $\mathcal{P}_j$, but only of $\mathcal{B}_j$.

We first present the following lemma:

Lemma 6 *For all* $1 \le j \le \gamma$,

$$\Pr\{\mathcal{D}_{j-1}|\mathcal{B}_{j-1}\} = \prod_{k=j}^{\gamma} \Phi(\mathcal{C}_k) = \Phi(\mathcal{C}_j)\Pr\{\mathcal{D}_j|\mathcal{B}_j\}.$$

Proof. For all $1 \le j \le \gamma$, as both $\mathcal{V} = \mathcal{P}_\gamma = \{\mathcal{D}_{j-1}, \mathcal{P}_{j-1}\}$ and $\mathcal{P}_{j-1}$ are parentless, from Theorem 2,

$$\begin{aligned}\Pr\{\mathcal{P}_\gamma\} = \prod_{k=0}^{\gamma} \Phi(\mathcal{C}_k) &= \Pr\{\mathcal{D}_{j-1}, \mathcal{P}_{j-1}\} = \Pr\{\mathcal{D}_{j-1}|\mathcal{P}_{j-1}\}\Pr\{\mathcal{P}_{j-1}\} \\ &= \Pr\{\mathcal{D}_{j-1}|\mathcal{B}_{j-1}\}\prod_{k=0}^{j-1} \Phi(\mathcal{C}_k).\end{aligned}$$

Assuming all $\Pr\{V|\mathcal{H}_V\} > 0$,

$$\Pr\{\mathcal{D}_{j-1}|\mathcal{B}_{j-1}\} = \prod_{k=j}^{\gamma} \Phi(\mathcal{C}_k) = \Phi(\mathcal{C}_j)\prod_{k=j+1}^{\gamma} \Phi(\mathcal{C}_k) = \Phi(\mathcal{C}_j)\Pr\{\mathcal{D}_j|\mathcal{B}_j\}.$$

■

We define the following notation: For all $0 \leq j \leq \gamma - 1$, by Lemma 4,

$$\Lambda\left(\mathcal{B}_j\right) = \Pr\left\{\left[\mathcal{D}_j\right]|\mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{B}_j} = \sum_{\mathcal{D}_j} \Pr\left\{\mathcal{D}_j|\mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{D}_j \cup \mathcal{B}_j}. \tag{11}$$

Since $\mathcal{D}_\gamma = \varnothing$, we also define $\Lambda\left(\mathcal{B}_\gamma\right) = \boldsymbol{I}_{\mathcal{B}_\gamma}$.

Although we write $\Lambda\left(\mathcal{B}_j\right)$, the scope of $\Lambda\left(\mathcal{B}_j\right)$ may not be the whole $\mathcal{B}_j$, because $\mathcal{B}_j$ may have more variables than the minimal set needed to separate $\mathcal{A}_j$ and $\mathcal{D}_j$. For example, in the BN $\mathbb{A}$, while $\mathcal{B}_4 = \{F, H, I\}$, we only need $\{H, I\}$ for the study of $\mathcal{D}_4 = \{G, J, K, L\}$.

Theorem 7 *For all* $1 \leq j \leq \gamma$,

$$\Lambda\left(\mathcal{B}_{j-1}\right) = \sum_{\mathcal{C}_j} \phi\left(\mathcal{C}_j\right) \Lambda\left(\mathcal{B}_j\right).$$

Proof. From Equation (6), $\mathcal{D}_{j-1} = \{\mathcal{C}_j, \mathcal{D}_j\}$. From Equation (4), $\mathcal{B}_{j-1} \cup \mathcal{C}_j = V_j \cup \mathcal{B}_j$. Also, if $V_j \neq \varnothing$, then its cohort $\mathcal{C}_j$ must include all its bottom children, or $V_j \in \mathcal{H}_{\mathcal{C}_j} \subseteq \mathcal{B}_{j-1}$. Thus,

$$\begin{aligned}\mathcal{B}_{j-1} \cup \mathcal{D}_{j-1} &= \mathcal{B}_{j-1} \cup \left\{\mathcal{C}_j \cup \mathcal{H}_{\mathcal{C}_j}\right\} \cup \left\{\mathcal{C}_j \cup \mathcal{D}_j\right\} \\ &= V_j \cup \mathcal{B}_j \cup \mathcal{H}_{\mathcal{C}_j} \cup \left\{\mathcal{C}_j \cup \mathcal{D}_j\right\} = \mathcal{B}_j \cup \mathcal{H}_{\mathcal{C}_j} \cup \mathcal{C}_j \cup \mathcal{D}_j.\end{aligned}$$

From Lemma 6,

$$\Lambda\left(\mathcal{B}_{j-1}\right) = \sum_{\mathcal{D}_{j-1}} \Pr\left\{\mathcal{D}_{j-1}|\mathcal{B}_{j-1}\right\} \boldsymbol{I}_{\mathcal{B}_{j-1} \cup \mathcal{D}_{j-1}} = \sum_{\{\mathcal{D}_j, \mathcal{C}_j\}} \Phi\left(\mathcal{C}_j\right) \boldsymbol{I}_{\mathcal{C}_j \cup \mathcal{H}_{\mathcal{C}_j}} \Pr\left\{\mathcal{D}_j|\mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{D}_j \cup \mathcal{B}_j}.$$

Because the scope of $\phi\left(\mathcal{C}_j\right)$ (which is $\left\{\mathcal{C}_j, \mathcal{H}_{\mathcal{C}_j}\right\} \subseteq \left\{\mathcal{C}_j, \mathcal{P}_{j-1}\right\} = \mathcal{P}_j$) is not in $\mathcal{D}_j$, from Equation (6),

$$\Lambda\left(\mathcal{B}_{j-1}\right) = \sum_{\mathcal{C}_j} \phi\left(\mathcal{C}_j\right) \sum_{\mathcal{D}_j} \Pr\left\{\mathcal{D}_j|\mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{D}_j \cup \mathcal{B}_j} = \sum_{\mathcal{C}_j} \phi\left(\mathcal{C}_j\right) \Lambda\left(\mathcal{B}_j\right).$$

■

Suppose there is a value of ω $(1 < \omega \leq \gamma)$ such that $\Lambda\left(\mathcal{B}_\omega\right)$ can be calculated. Especially, let β be the last time an evidence variable is recruited into $\mathcal{P}$. Then for all $\beta \leq j \leq \gamma$, $\mathcal{D}_j$ has no evidence. Hence

$$\Lambda\left(\mathcal{B}_j\right) = \boldsymbol{I}_{\mathcal{B}_j} \text{ for all } \beta \leq j \leq \gamma. \tag{12}$$

Starting with $\Lambda(\mathcal{B}_\omega)$, we can calculate $\Lambda(\mathcal{B}_j)$ for all $0 \leq j < \omega$ recursively by the above lemma.

For the BN $\mathbb{A}$, with $\mathcal{E} = \{H = h, K = k\}$. Thus $\beta = 6$ and $\Lambda(\mathcal{B}_8) = \Lambda(\mathcal{B}_7) = \Lambda(\mathcal{B}_6) = \boldsymbol{I}_K$.

1. Because $\mathcal{B}_6$ has cohort $\mathcal{C}_6 = \{J, K, G\}$:

$$\begin{aligned}\Lambda(\mathcal{B}_5) &= \sum_{\mathcal{C}_6} \phi(\mathcal{C}_6)\Lambda(\mathcal{B}_6) \\ &= \Pr\{k|h, I\} \sum_G \left(\Pr\{G\} \sum_J \Pr\{J|G, h\} \right) \boldsymbol{I}_K = \Pr\{k|I, h\}.\end{aligned}$$

2. Because $\mathcal{B}_5$ has cohort $\mathcal{C}_5 = \varnothing$:

$$\Lambda(\mathcal{B}_4) = \Lambda(\mathcal{B}_5) = \Pr\{k|I, h\}.$$

3. Because $\mathcal{B}_4$ has cohort $\mathcal{C}_4 = I$:

$$\begin{aligned}\Lambda(\mathcal{B}_3) &= \sum_{\mathcal{C}_4} \phi(\mathcal{C}_4)\Lambda(\mathcal{B}_4) = \sum_I \Pr\{I|D, F\} \Pr\{k|I, h\} \\ &= \sum_I \Pr\{k, I|D, F, h\} = \Pr\{k|D, F, h\}.\end{aligned}$$

4. Because $\mathcal{B}_3$ has cohort $\mathcal{C}_3 = H = h$:

$$\Lambda(\mathcal{B}_2) = \phi(\mathcal{C}_3)\Lambda(\mathcal{B}_3) = \Pr\{h|C, D\} \Pr\{k|D, F, h\} = \Pr\{h, k|C, D, F\}.$$

5. Because $\mathcal{B}_2$ has cohort $\mathcal{C}_2 = \varnothing$:

$$\Lambda(\mathcal{B}_1) = \Lambda(\mathcal{B}_2) = \Pr\{h, k|C, D, F\}.$$

6. Because $\mathcal{B}_1$ has cohort $\mathcal{C}_1 = \{C, D, F\}$:

$$\begin{aligned}\Lambda(\mathcal{B}_0) &= \sum_{\mathcal{C}_1} \phi(\mathcal{C}_1)\Lambda(\mathcal{B}_1) \\ &= \sum_{\{C,D,F\}} \Pr\{C|A, B\} \Pr\{D|A, B\} \Pr\{F|A, B\} \Pr\{h, k|C, D, F\} \\ &= \Pr\{h, k|A, B\}.\end{aligned}$$

4.4 The Posterior Marginal Distributions

Combining the downward and upward passes yields:

Theorem 8 *For all* $0 \leq j \leq \gamma$,

$$\Pr\left\{\mathcal{B}_j, [\mathcal{E}\backslash\mathcal{B}_j]\right\} \boldsymbol{I}_{\mathcal{B}_j} = \Pi\left(\mathcal{B}_j\right) \Lambda\left(\mathcal{B}_j\right).$$

Proof. By Lemma 4, because the event $[\mathcal{E}\backslash\mathcal{B}_j]$ is the same as the event $[\mathcal{V}\backslash\mathcal{B}_j] = [\mathcal{A}_j \cup \mathcal{D}_j]$,

$$\begin{aligned}\Pr\left\{\mathcal{B}_j, [\mathcal{E}\backslash\mathcal{B}_j]\right\} \boldsymbol{I}_{\mathcal{B}_j} &= \Pr\left\{\mathcal{B}_j, [\mathcal{A}_j \cup \mathcal{D}_j]\right\} \boldsymbol{I}_{\mathcal{B}_j} = \sum_{\{\mathcal{A}_j, \mathcal{D}_j\}} \Pr\left\{\mathcal{B}_j, \mathcal{A}_j, \mathcal{D}_j\right\} \boldsymbol{I}_{\mathcal{B}_j \cup \mathcal{A}_j \cup \mathcal{D}_j} \\ &= \sum_{\{\mathcal{A}_j, \mathcal{D}_j\}} \Pr\left\{\mathcal{A}_j, \mathcal{B}_j\right\} \Pr\left\{\mathcal{D}_j | \mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{B}_j \cup \mathcal{A}_j \cup \mathcal{D}_j}.\end{aligned}$$

As $\mathcal{D}_j \cap \{\mathcal{A}_j, \mathcal{B}_j\} = \varnothing$,

$$\begin{aligned}\Pr\left\{\mathcal{B}_j, [\mathcal{E}\backslash\mathcal{B}_j]\right\} \boldsymbol{I}_{\mathcal{B}_j} &= \sum_{\mathcal{A}_j} \Pr\left\{\mathcal{A}_j, \mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{A}_j \cup \mathcal{B}_j} \left(\sum_{\mathcal{D}_j} \Pr\left\{\mathcal{D}_j | \mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{D}_j \cup \mathcal{B}_j} \right) \\ &= \sum_{\mathcal{A}_j} \Pr\left\{\mathcal{A}_j, \mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{A}_j \cup \mathcal{B}_j} \Lambda\left(\mathcal{B}_j\right).\end{aligned}$$

As $\mathcal{B}_j \cap \mathcal{A}_j = \varnothing$,

$$\Pr\left\{\mathcal{B}_j, [\mathcal{E}\backslash\mathcal{B}_j]\right\} \boldsymbol{I}_{\mathcal{B}_j} = \Lambda\left(\mathcal{B}_j\right) \sum_{\mathcal{A}_j} \Pr\left\{\mathcal{A}_j, \mathcal{B}_j\right\} \boldsymbol{I}_{\mathcal{A}_j \cup \mathcal{B}_j} = \Pi\left(\mathcal{B}_j\right) \Lambda\left(\mathcal{B}_j\right).$$

■

Corollary 9 *For node* $Q \in \mathcal{B}_j$ *where* $0 \leq j \leq \gamma$,

$$\Pr\left\{Q, [\mathcal{E}\backslash Q]\right\} \boldsymbol{I}_Q = \sum_{\mathcal{B}_j \backslash Q} \Pi\left(\mathcal{B}_j\right) \Lambda\left(\mathcal{B}_j\right).$$

If $Q \notin \mathcal{E}$,

$$\Pr\left\{Q, [\mathcal{E}]\right\} = \sum_{\mathcal{B}_j \backslash Q} \Pi\left(\mathcal{B}_j\right) \Lambda\left(\mathcal{B}_j\right).$$

Proof. For node $Q \in \mathcal{B}_j$,

$$\begin{aligned}\Pr\{Q, [\mathcal{E}\backslash Q]\}\, \boldsymbol{I}_Q &= \sum_{\mathcal{E}\backslash Q} \Pr\{\mathcal{E}\}\, \boldsymbol{I}_{\mathcal{E}} = \sum_{\mathcal{B}_j\backslash Q} \boldsymbol{I}_{\mathcal{B}_j} \sum_{\mathcal{E}\backslash\mathcal{B}_j} \Pr\{\mathcal{B}_j, \mathcal{E}\backslash\mathcal{B}_j\}\, \boldsymbol{I}_{\mathcal{E}\backslash\mathcal{B}_j} \\ &= \sum_{\mathcal{B}_j\backslash Q} \boldsymbol{I}_{\mathcal{B}_j} \Pr\{\mathcal{B}_j, [\mathcal{E}\backslash\mathcal{B}_j]\} = \sum_{\mathcal{B}_j\backslash Q} \Pi(\mathcal{B}_j)\, \Lambda(\mathcal{B}_j).\end{aligned}$$

■

Recall that β is the last time an evidence is recruited into $\mathcal{P}$, if we are looking for the "post-evidence" $\Pr\{Q, [\mathcal{E}\backslash Q]\}\, \boldsymbol{I}_Q$ where $Q \in \mathcal{B}_j$ and $\beta \leq j \leq \gamma$, then due to Equation (12), we can find them by the downward pass alone as $\sum_{\mathcal{B}_j\backslash Q} \Pi(\mathcal{B}_j)\, \boldsymbol{I}_{\mathcal{B}_j} = \sum_{\mathcal{B}_j\backslash Q} \Pi(\mathcal{B}_j)$.

For the BN $\mathbb{A}$ with $\mathcal{E} = \{H = h, K = k\}$, by the downward pass alone we already have

$$\begin{aligned}\Pr\{\mathcal{B}_8, [\mathcal{E}\backslash\mathcal{B}_8]\}\, \boldsymbol{I}_{\mathcal{B}_8} &= \Pi(\mathcal{B}_8) = \Pr\{J, L, k, h\}; \\ \Pr\{\mathcal{B}_7, [\mathcal{E}\backslash\mathcal{B}_7]\}\, \boldsymbol{I}_{\mathcal{B}_7} &= \Pi(\mathcal{B}_7) = \Pr\{I, J, k, h\}; \\ \Pr\{\mathcal{B}_6, [\mathcal{E}\backslash\mathcal{B}_6]\}\, \boldsymbol{I}_{\mathcal{B}_6} &= \Pi(\mathcal{B}_6) = \Pr\{I, J, G, k, h\}.\end{aligned}$$

Now with Theorem 8,

1. $\Pi(\mathcal{B}_5)\, \Lambda(\mathcal{B}_5) = \Pr\{I, h\} \Pr\{k|I, h\} = \Pr\{I, h, k\}$.
2. $\Pi(\mathcal{B}_4)\, \Lambda(\mathcal{B}_4) = \Pr\{F, I, h\} \Pr\{k|I, h\} = \Pr\{F, I, h, k\}$.
3. $\Pi(\mathcal{B}_3)\, \Lambda(\mathcal{B}_3) = \Pr\{D, F, h\} \Pr\{k|D, F, h\} = \Pr\{D, F, h, k\}$.
4. $\Pi(\mathcal{B}_2)\, \Lambda(\mathcal{B}_2) = \Pr\{C, D, F\} \Pr\{h, k|C, D, F\} = \Pr\{C, D, F, h, k\}$.
5. $\Pi(\mathcal{B}_1)\, \Lambda(\mathcal{B}_1) = \Pr\{B, C, D, F\} \Pr\{h, k|C, D, F\} = \Pr\{B, C, D, F, h, k\}$.
6. $\Pi(\mathcal{B}_0)\, \Lambda(\mathcal{B}_0) = \Pr\{A, B\} \Pr\{h, k|A, B\} = \Pr\{A, B, h, k\}$.

A variable $V \in \mathcal{V}$ may appear in more than one borders. For example, variable I appears in $\mathcal{B}_4$, $\mathcal{B}_5$, $\mathcal{B}_6$ and $\mathcal{B}_7$. We obtain the same result regardless which of these borders we choose to marginalize.

Border algorithm is applicable to all BNs. In the next section, we study a special kind of BNs called the polytrees, and present in details, with some modifications and within the border algorithm framework, the "revised polytree algorithm" by Peot & Shachter (1991). This is an important section, because we will show later that, with the help of the border algorithm, any BN can be modified to become a polytree.

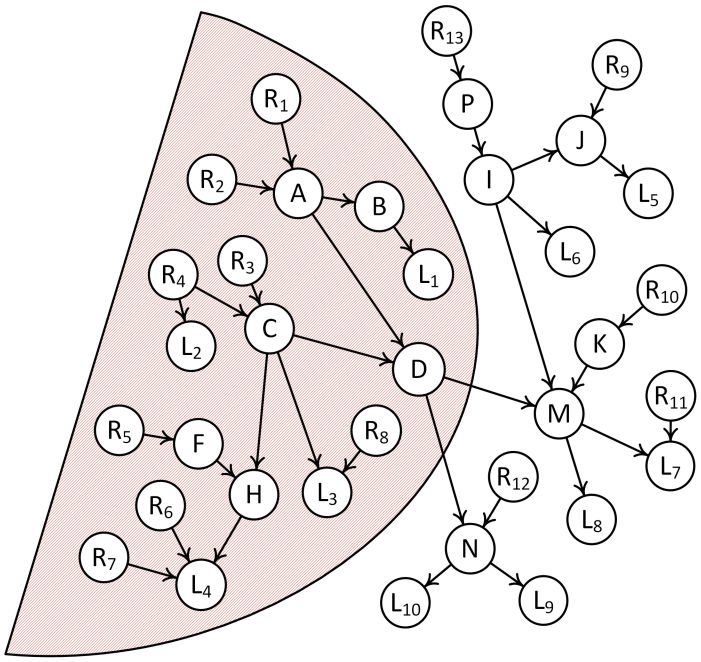

Figure 4: The Polytree $\mathbb{B}$

5 The Revised Polytree Algorithm

A polytree is a BN which is "singly connected;" that is, there is only one *undirected* path connecting any two nodes. (From now on, "path" means "undirected path.") The BN $\mathbb{A}$ in Figure 1 is not a polytree because there are 2 paths from A to H, namely $A - C - H$ and $A - D - H$.

In other words, while we assume all BNs are acyclic (that is, they have no directed cycles), a polytree also does not have any undirected cycle (or "loop"). The BN $\mathbb{A}$ has loop $A - C - H - D - A$.

As an illustration, we will use the polytree as shown in Figure 4, which we will refer to as the Polytree $\mathbb{B}$.

5.1 To Find the Path Connecting Two Nodes

Here we present a method to identify the unique path connecting any two nodes in a polytree.

We strategically designate some nodes as "hubs," and pre-load the unique

path connecting each pair of hubs, excluding the hubs themselves. The bigger the network, the more hubs we need. For the Polytree $\mathbb{B}$, let us pick nodes J and H as hubs, connected by path $I - M - D - C$.

For each node, we also pre-load the path from it to its nearest hub. For node P the path is $P - I - J$; for node A the path is $A - D - C - H$.

To find the path from node X to node Y:

1. Form the (possibly cyclic) path from X to the hub nearest to X, then to the hub nearest to Y, then to Y. For nodes P and A, this path is $(P - I - J) - (I - M - D - C) - (H - C - D - A)$.

2. Replace the largest loop around each hub with its furthest node. With the above path, replace loop $I - J - I$ with node I, and loop $D - C - H - C - D$ with node D, resulting in path $P - I - M - D - A$ connecting nodes P and A.

5.2 The Decompositions by Nodes

Node V in a polytree decomposes the polytree into two parts:

1. $\mathcal{P}_V$, the parentless set including V and all the nodes that are connected to V "from above," via its parents $\mathcal{H}_V$. (In Figure 4, $\mathcal{P}_D$ is in the shaded region.) $\mathcal{P}_V$ has border V and the top part $\mathcal{A}_V = \mathcal{P}_V \backslash V$. Consistent with Definition (9), we define

$$\Pi(V) = \Pr\{V, [\mathcal{A}_V]\} \boldsymbol{I}_V = \sum_{\mathcal{A}_V} \Pr\{\mathcal{P}_V\} \boldsymbol{I}_{\mathcal{P}_V}. \tag{13}$$

2. $\mathcal{D}_V$, the bottom set of $\mathcal{P}_V$, in which all nodes are connected to V "from below," via its children $\mathcal{L}_V$. Consistent with Definition (11), we define

$$\Lambda(V) = \Pr\{[\mathcal{D}_V] | V\} \boldsymbol{I}_V = \sum_{\mathcal{D}_V} \Pr\{\mathcal{D}_V | V\} \boldsymbol{I}_{V \cup \mathcal{D}_V}. \tag{14}$$

5.3 The Decompositions by Edges

So far, we focused on the nodes in a BN. Let us now consider a typical edge $X \to Y$. While node X decomposes the polytree into $\mathcal{P}_X$ and $\mathcal{D}_X$, edge $X \to Y$ also decomposes it into two parts:

1. $\mathcal{T}_{X\to Y}$, the parentless set of nodes on the parent side of edge $X \to Y$, having border X. Not only does $\mathcal{T}_{X\to Y}$ include $\mathcal{P}_X$, but also all the nodes that connect to X from below, except those via Y. In the Polytree $\mathbb{B}$, in addition to $\mathcal{P}_D$, $\mathcal{T}_{D\to M}$ also includes $\{N, R_{12}, L_9, L_{10}\}$.

   Consistent with Definition (9), we define the *downward message* about the evidences in $\mathcal{T}_{X\to Y}$ that node X can send to its child Y as

$$\Pi_Y(X) = \Pr\{X, [\mathcal{T}_{X\to Y}\backslash X]\}\, \boldsymbol{I}_X = \sum_{\mathcal{T}_{X\to Y}\backslash X} \Pr\{\mathcal{T}_{X\to Y}\}\, \boldsymbol{I}_{\mathcal{T}_{X\to Y}}. \quad (15)$$

2. The bottom set $\mathcal{U}_{X\to Y} = \mathcal{V}\backslash\mathcal{T}_{X\to Y}$, on the child side of edge $X \to Y$, separated from $\mathcal{T}_{X\to Y}\backslash X$ by X. Not only does $\mathcal{U}_{X\to Y}$ include $\{Y, \mathcal{D}_Y\}$, but also all the nodes that connect to Y from above, except those via X. Hence,

$$\mathcal{U}_{X\to Y} = \left\{Y, \mathcal{D}_Y, \cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}\right\}. \quad (16)$$

   On the other hand,

$$\mathcal{T}_{X\to Y} = \left\{\mathcal{P}_X, \cup_{V\in\mathcal{L}_X\backslash Y}\mathcal{U}_{X\to V}\right\}. \quad (17)$$

   Consistent with Definition (11), we define the *upward message* about the evidences in $\mathcal{U}_{X\to Y}$ that node Y can send to its parent X as

$$\Lambda_Y(X) = \Pr\{[\mathcal{U}_{X\to Y}]\,|X\}\, \boldsymbol{I}_X = \sum_{\mathcal{U}_{X\to Y}} \Pr\{\mathcal{U}_{X\to Y}|X\}\, \boldsymbol{I}_{X\cup\mathcal{U}_{X\to Y}}. \quad (18)$$

We will often use the following two properties related to edge $X \to Y$ in a polytree:

1. Because two distinct parents Z and T of X have X as a common child, the two parentless sets $\mathcal{T}_{Z\to X}$ and $\mathcal{T}_{T\to X}$ must be disjoint and independent (otherwise, there are two paths from their common member to Y, via the two parents). Thus,

$$\Pr\{\cup_{V\in\mathcal{H}_X}\mathcal{T}_{V\to X}\} = \prod_{V\in\mathcal{H}_X}\Pr\{\mathcal{T}_{V\to X}\}.$$

2. Because two distinct children Y and W of X have X as a common parent, the two sets $\mathcal{U}_{X\to Y}$ and $\mathcal{U}_{X\to W}$ must be disjoint and independent given X (otherwise, there are two paths from X to their common member, via the two children). Thus,

$$\Pr\left\{\cup_{V\in\mathcal{L}_X}\mathcal{U}_{X\to V}|X\right\}=\prod_{V\in\mathcal{L}_X}\Pr\left\{\mathcal{U}_{X\to V}|X\right\}.$$

5.4 The Message Propagations in the Polytrees

We now present the lemmas about the relationships among $\Pi(X)$, $\Lambda(X)$, $\Pi_Y(X)$ and $\Lambda_Y(X)$, which are known in the literature, but are now proven within the border algorithm framework:

Lemma 10 *For edge $X\to Y$ in a polytree,*

$$\Pi_Y(X)=\Pi(X)\prod_{V\in\mathcal{L}_X\setminus Y}\Lambda_V(X).$$

Proof. Consider the parentless $\mathcal{P}_X$, having border X. As in Equation (17), recruiting the childless $\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}$ without promotion results in the parentless $\mathcal{T}_{X\to Y}$ with border $X\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}$. From Theorem 5,

$$\begin{aligned}\Pi\left(X\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}\right)&=\Pi(X)\Pr\left\{\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}|X\right\}\boldsymbol{I}_{X\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}}\\&=\Pi(X)\prod_{V\in\mathcal{L}_X\setminus Y}\Pr\left\{\mathcal{U}_{X\to V}|X\right\}\boldsymbol{I}_{X\cup\mathcal{U}_{X\to V}}.\end{aligned}$$

Now use Rule 1 to promote the childless $\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}$ without cohort, resulting in the parentless $\mathcal{T}_{X\to Y}$ having border X. From Theorem 5,

$$\begin{aligned}\Pi_Y(X)&=\sum_{\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}}\Pi\left(X\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}\right)\\&=\sum_{\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}}\Pi(X)\prod_{V\in\mathcal{L}_X\setminus Y}\Pr\left\{\mathcal{U}_{X\to V}|X\right\}\boldsymbol{I}_{X\cup\mathcal{U}_{X\to V}}.\end{aligned}$$

As the scope of $\Pi(X)$ (which is X) is not in $\cup_{V\in\mathcal{L}_X\setminus Y}\mathcal{U}_{X\to V}$,

$$\Pi_Y(X)=\Pi(X)\prod_{V\in\mathcal{L}_X\setminus Y}\left(\sum_{\mathcal{U}_{X\to V}}\Pr\left\{\mathcal{U}_{X\to V}|X\right\}\boldsymbol{I}_{X\cup\mathcal{U}_{X\to V}}\right),$$

which is the lemma by Definition (18). This is similar to Equation (4.45) in Pearl, 1988. ■

Lemma 11 *For edge $X \to Y$ in a polytree,*

$$\Lambda_Y(X) = \sum_Y \Lambda(Y) \sum_{\mathcal{H}_Y \backslash X} \Pr_r\{Y|\mathcal{H}_Y\} \prod_{V \in \mathcal{H}_Y \backslash X} \Pi_Y(V).$$

Proof. Consider the parentless $\mathcal{T}_{X\to Y}$, having border X and the bottom set $\mathcal{U}_{X\to Y}$. As shown in Equation (16), recruiting $Y \cup_{V\in\mathcal{H}_Y\backslash X} \mathcal{T}_{V\to Y}$ without promotion results in the parentless $\mathcal{P}_Y$ with bottom part $\mathcal{D}_Y$. The reduced cohort table is

$$\begin{aligned}&\Pr\left\{Y \cup_{V\in\mathcal{H}_Y\backslash X} \mathcal{T}_{V\to Y}|X\right\} \boldsymbol{I}_{X\cup Y\cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}} \\ &= \Pr\left\{Y|\cup_{V\in\mathcal{H}_Y\backslash X} \mathcal{T}_{V\to Y}, X\right\} \Pr\left\{\cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}|X\right\} \boldsymbol{I}_{X\cup Y\cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}}\end{aligned}$$

Because $\mathcal{H}_Y \subseteq X \cup_{V\in\mathcal{H}_Y\backslash X} \mathcal{T}_{V\to Y}$, the reduced cohort table becomes

$$\begin{aligned}&\Pr\{Y|\mathcal{H}_Y\} \boldsymbol{I}_{Y\cup\mathcal{H}_Y} \Pr\left\{\cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}\right\} \boldsymbol{I}_{\cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}} \\ &= \Pr_r\{Y|\mathcal{H}_Y\} \prod_{V\in\mathcal{H}_Y\backslash X} \Pr\{\mathcal{T}_{V\to Y}\} \boldsymbol{I}_{\mathcal{T}_{V\to Y}}.\end{aligned}$$

From Theorem 7, as $Y \notin \cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}$,

$$\begin{aligned}\Lambda_Y(X) &= \sum_{Y\cup_{V\in\mathcal{H}_Y\backslash X}\mathcal{T}_{V\to Y}} \left(\Pr_r\{Y|\mathcal{H}_Y\} \prod_{V\in\mathcal{H}_Y\backslash X} \Pr\{\mathcal{T}_{V\to Y}\} \boldsymbol{I}_{\mathcal{T}_{V\to Y}}\right) \Lambda(Y) \\ &= \sum_Y \Lambda(Y) \sum_{\mathcal{H}_Y\backslash X} \Pr_r\{Y|\mathcal{H}_Y\} \prod_{V\in\mathcal{H}_Y\backslash X} \left(\sum_{\mathcal{T}_{V\to Y}\backslash V} \Pr\{\mathcal{T}_{V\to Y}\} \boldsymbol{I}_{\mathcal{T}_{V\to Y}}\right),\end{aligned}$$

which is the lemma by Definition (15). This is similar to Equation (4.44) in Pearl (1988). ■

The values of $\Pi(\cdot)$ and $\Lambda(\cdot)$ can be calculated from the messages as in the following lemmas:

Lemma 12 *For node X in a polytree,*

$$\Pi(X) = \sum_{\mathcal{H}_X} \Pr_r\{X|\mathcal{H}_X\} \prod_{V\in\mathcal{H}_X} \Pi_X(V).$$

Proof. Consider the parentless $\cup_{V\in\mathcal{H}_X}\mathcal{T}_{V\to X}$ having border $\mathcal{H}_X$ and the top part $\cup_{V\in\mathcal{H}_X}\{\mathcal{T}_{V\to X}\backslash V\}$. By Definition (9),

$$\begin{aligned}\Pi(\mathcal{H}_X) &= \sum_{\cup_{V\in\mathcal{H}_X}\{\mathcal{T}_{V\to X}\backslash V\}} \Pr\{\cup_{V\in\mathcal{H}_X}\mathcal{T}_{V\to X}\}\, \boldsymbol{I}_{\cup_{V\in\mathcal{H}_X}\mathcal{T}_{V\to X}} \\ &= \prod_{V\in\mathcal{H}_X}\left(\sum_{\mathcal{T}_{V\to X}\backslash V}\Pr\{\mathcal{T}_{V\to X}\}\,\boldsymbol{I}_{\mathcal{T}_{V\to X}}\right) = \prod_{V\in\mathcal{H}_X}\Pi_X(V).\end{aligned}$$

Now recruit X, resulting in the parentless $\mathcal{P}_X = X\cup_{V\in\mathcal{H}_X}\mathcal{T}_{V\to X}$. Then use Rule 1 to promote $\mathcal{H}_X$ without cohort, leaving X as the border of $\mathcal{P}_X$. From Theorem 5,

$$\Pi(X) = \sum_{\mathcal{H}_X}\Pr_r\{X|\mathcal{H}_X\}\,\Pi(\mathcal{H}_X) = \sum_{\mathcal{H}_X}\Pr_r\{X|\mathcal{H}_X\}\prod_{V\in\mathcal{H}_X}\Pi_X(V),$$

hence the lemma, which is similar to Equation (4.38) in Pearl (1988). ■

Lemma 13 *For node X in a polytree,*

$$\Lambda(X) = \prod_{V\in\mathcal{L}_X}\Lambda_V(X).$$

Proof. Consider the parentless $\mathcal{P}_X$, having border X. Recruit the rest of the network $\mathcal{D}_X = \cup_{V\in\mathcal{L}_X}\mathcal{U}_{X\to V}$ without promotion, resulting in the bottom part $\varnothing$ with $\Lambda(\varnothing)=1$. From Theorem 7,

$$\begin{aligned}\Lambda(X) &= \sum_{\cup_{V\in\mathcal{L}_X}\mathcal{U}_{X\to V}}\Pr\{\cup_{V\in\mathcal{L}_X}\mathcal{U}_{X\to V}|X\}\,\boldsymbol{I}_{X\cup_{V\in\mathcal{L}_X}\mathcal{U}_{X\to V}}\Lambda(\varnothing) \\ &= \prod_{V\in\mathcal{L}_X}\sum_{\mathcal{U}_{X\to V}}\Pr\{\mathcal{U}_{X\to V}|X\}\,\boldsymbol{I}_{X\cup\mathcal{U}_{X\to V}},\end{aligned}$$

hence the lemma by Definition (18), which is similar to Equation (4.35) in Pearl (1988). ■

Our dealing with the evidences here is slightly different to that in Peot & Shachter (1991, pp. 308-309). While we attach the indicator column $\boldsymbol{I}_V$ to all $\Pi(V)$, $\Pi_Y(V)$, $\Lambda(V)$ and $\Lambda_Y(V)$, they call it the "local evidence" $\Lambda_V(V)$ and attach it to $\Lambda(V)$ only.

Combining Lemmas 10 and 12 yields:

Theorem 14 *If node X has received the messages from all members of $\mathcal{H}_X \cup \{\mathcal{L}_X \backslash Y\}$, then it can send a downward message to its child Y as*

$$\Pi_Y(X) = \left(\prod_{V \in \mathcal{L}_X \backslash Y} \Lambda_V(X) \right) \sum_{\mathcal{H}_X} \left(\Pr_r\{X|\mathcal{H}_X\} \prod_{V \in \mathcal{H}_X} \Pi_X(V) \right).$$

Combining Lemmas 11 and 13 yields:

Theorem 15 *If node X has received the messages from all members of $\{\mathcal{H}_X \backslash H\} \cup \mathcal{L}_X$, then it can send an upward message to is parent H as*

$$\Lambda_X(H) = \sum_X \left(\prod_{V \in \mathcal{L}_X} \Lambda_V(X) \right) \sum_{\mathcal{H}_X \backslash H} \left(\Pr_r\{X|\mathcal{H}_X\} \prod_{V \in \mathcal{H}_X \backslash H} \Pi_X(V) \right).$$

5.5 The Collection Phase

In the first phase of the revised polytree algorithm, which is known as the "collection phase," a randomly selected node P is designated as a "pivot," and the messages are passed (or "propagated") to P.

5.5.1 The Evidential Cores

The "evidential core" (or EC) of a polytree is the smallest sub-polytree which contains all the evidences. In other words, it comprises of the evidence set $\mathcal{E}$ and all the nodes and edges on the path connecting every pair of the evidence nodes. Corresponding to a particular evidence set, the EC is a unique.

Figure 5 shows the EC in the Polytree $\mathbb{B}$, corresponding to $\mathcal{E} = \{B, C, K, L_4\}$, not including the nodes or edges in dash, such as node L_1. We call it the EC $\mathbb{B}$.

A node in a polytree is said to be "inside" if it is in the EC; otherwise it is "outside." Likewise, an edge $X \to Y$ is "inside" if both X and Y are inside; otherwise it is "outside." A path is "inside" if all its component edges are.

It is important to observe that, for an outside edge $X \to Y$, the whole EC can only be either on its parent side $\mathcal{T}_{X \to Y}$, or its child side $\mathcal{U}_{X \to Y}$, not both. Otherwise, being connected, the EC must include edge $X \to Y$ on the path connecting the evidences on both sides.

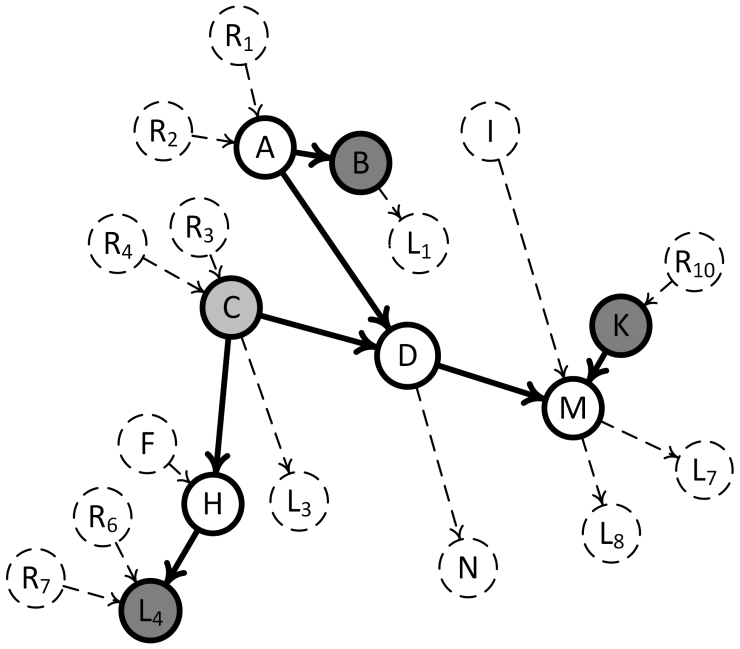

Figure 5: The EC $\mathbb{B}$ in the Polytree $\mathbb{B}$, corresponding to $\mathcal{E} = \{B, C, K, L_4\}$

5.5.2 The Boundary Conditions

We want to use Theorems 14 and 15 for calculating the messages along the inside edges only. However, this requires the knowledge of the boundary conditions; that is, of the messages from the outside neighbors to the inside nodes, such as from those in dash in Figure 5. So we need the following theorem:

Theorem 16 *Consider an inside node X of an EC.*
(a) If $V \in \mathcal{H}_X$ is outside, then $\Pi_X(V) = \Pr\{V\}$.
(b) If $V \in \mathcal{L}_X$ is outside, then $\Lambda_V(X) = \boldsymbol{I}_X$.

Proof. (a) Because the inside X is on the child side of edge $V \to X$, so is the whole EC, and thus $\mathcal{T}_{V\to X}$ has no evidence, or $\boldsymbol{I}_{\mathcal{T}_{V\to X}} = 1$. Hence, by Definition (15), $\Pi_X(V) = \Pr\{V\}$.

(b) If $V \in \mathcal{L}_X$ is outside, then because the inside X is on the parent side of edge $X \to V$, so is the whole EC; hence $\mathcal{U}_{X\to V}$ has no evidence. By Definition (18), $\Lambda_V(X) = \boldsymbol{I}_X$. ■

So, assuming that all prior probabilities $\Pr\{V\}$ are pre-loaded, Theorems 14 and 15 can be used without the need to calculate the messages from the outside neighbors.

5.5.3 The Message Initializations

Like all polytrees, an EC has its own roots and leaves. For example, the sets of roots and leaves of the EC $\mathbb{B}$ are $\{K, A, C\}$ and $\{B, L_4, M\}$, respectively.

By definition of the EC, every path in it must end with an evidence node, although there are evidence nodes not at the end of a path (such as nodes C in the EC $\mathbb{B}$). Also, an evidence node at the end of a path can only have one inside child or parent, otherwise the path does not end with it. If it has one child, then it is a root of the EC (such as nodes K in the EC $\mathbb{B}$); if one parent, it is a leaf (such as nodes B and L_4 in the EC $\mathbb{B}$).

We initialize the message propagations in an EC as follows:

Theorem 17 *Consider an EC.*
(a) For an evidence root R having one inside child C, $\Pi_C(R) = \Pi(R) = \Pr\{R\}\boldsymbol{I}_R$.
(b) For an evidence leaf L, $\Lambda(L) = \boldsymbol{I}_L$.

Proof. (a) Because $\mathcal{A}_R = \mathcal{P}_R \backslash R$ is non-evidential, $\Pi(R) = \Pr\{R\} \boldsymbol{I}_R$ by Definition (13). Also, because R has only one inside child C, $\Lambda_V(R) = \boldsymbol{I}_R$ for all $V \in \mathcal{L}_R \backslash C$. Thus $\Pi_C(R) = \Pi(R)$ from Lemma 10.

(b) For an evidence leaf L, $\mathcal{D}_L$ has no evidence and thus $\Lambda(L) = \boldsymbol{I}_L$ by Definition (14). ■

5.5.4 To the Pivot

In the collection phase, starting with an evidence root or leaf in the EC, we send the messages from all the inside nodes to a randomly selected inside pivot node P. At any time, if P has not received a message along an edge to or from it, we trace along the possibly many inside paths leading to this edge, and must see either an inside node that is ready to do so, or an evidence node at the end of an inside path that can send its message to P by Theorem 17.

We may need both Theorems 14 and 15 to send the messages along a path, because the directions may change (for example, along path $B \leftarrow A \rightarrow D$). In the EC $\mathbb{B}$ in Figure 5, if we use node D as the pivot, then the collection phase includes the messages sent along paths $B \leftarrow A \rightarrow D$, $L_4 \leftarrow H \leftarrow C \rightarrow D$ and $K \rightarrow M \leftarrow D$.

5.6 The Distribution Phase

In the second phase of the revised polytree algorithm, which is known as the "distribution phase," the messages are passed from the pivot node P to a query variable.

5.6.1 The Informed Nodes

Once a node Q has received the messages from all its neighbors, we say it is "informed." At the end of the collection phase, the pivot node is the first informed node.

The posterior marginal probability of an informed node can now be obtained: With the messages from all its parents, we can use Lemma 12 to calculate $\Pi(Q)$; with the messages from all its children, we can use Lemma 13 to calculate $\Lambda(Q)$. $\Pr\{Q, [\mathcal{E} \backslash Q]\} \boldsymbol{I}_Q$ then can be calculated by Theorem 8 as $\Pi(Q)\Lambda(Q)$. Alternatively, we can use the following theorem:

Theorem 18 *With edge $Q \to Y$,*

$$\Pr\{Q, [\mathcal{E}\backslash Q]\}\, \boldsymbol{I}_Q = \Pi_Y(Q)\, \Lambda_Y(Q).$$

Proof. By Theorem 8, Lemmas 13 and 10,

$$\begin{aligned}\Pr\{Q, [\mathcal{E}\backslash Q]\}\, \boldsymbol{I}_Q &= \Pi(Q)\, \Lambda(Q) = \Pi(Q) \prod_{V \in \mathcal{L}_Q} \Lambda_V(Q) \\ &= \left(\Pi(Q) \prod_{V \in \mathcal{L}_Q \backslash Y} \Lambda_V(Q)\right) \Lambda_Y(Q) = \Pi_Y(Q)\, \Lambda_Y(Q).\end{aligned}$$

■

5.6.2 The Inside Query Variables

Let p be the number of paths connecting an inside node V with the rest of the network (that is, the number of its neighbors). If it is not the pivot, then only one of these paths leads to the pivot node P. For us to use Theorem 14 or 15 to send a message from V to P in the collection phase, V must have received the messages along all $p-1$ paths, except the one leading to P. Now, in the distribution phase, once it receives a message from the informed P, it becomes informed.

Let $\mathcal{J}$ be the set of all the informed nodes, which we call the "informed set." When we sent the messages from an informed node V to an uninformed node Q, not only Q, but also all the nodes along the path from V to Q become informed. In the EC $\mathbb{B}$ in Figure 5, the message propagations from the informed node D to node H also make node C informed. Thus, "spreading out" from a single pivot node, the informed set forms a connected sub-polytree, in that there is a single path connecting any two informed nodes and all nodes along that path are informed.

5.6.3 The Outside Query Variables

Starting from an inside pivot node, the informed set $\mathcal{J}$ does not have to cover the EC and can spread beyond it.

Consider now an outside uninformed node Q. The paths starting from Q to all nodes in $\mathcal{J}$ must go through a unique "gate" $T \in \mathcal{J}$; otherwise, because $\mathcal{J}$ is connected, there are more than one paths connecting Q with

an informed node via the different gates. Thus Q only needs the messages sent from T to it (in this direction) to become informed, as this has all the evidential information. All other messages sent to it along other paths are from the outside neighbors. In this manner, the informed set $\mathcal{J}$ spreads to Q.

For the Polytree $\mathbb{B}$ in Figure 5, suppose $\mathcal{J} = \{D, C, H\}$ and the outside node R_8 is a query variable. The gate from R_8 to $\mathcal{J}$ is C and the messages along path $C \leftarrow L_3 \rightarrow R_8$ are all that are needed to make R_8 informed.

Peot & Shachter's revised polytree algorithm has been very much neglected in the literature because not many practical BNs have this form. While they further suggested that their algorithm be applied to a general BN via "cutset conditioning," we continue this paper by presenting a novel method to convert any BN into a polytree.

6 The Border Polytrees (BPs)

In Section 2, we showed that the border algorithm "stretches" a BN into a border chain, which is a special form of the polytrees, in which each node has at most one parent and one child. For example, we stretched the BN $\mathbb{A}$ in Figure 1 into a border chain in Figure 3. We will now show how to convert any BN into a polytree.

6.1 Stage I: The Macro-node Polytrees

Our method involves two stages: In Stage I, we partition the BN to form a polytree.

A set of nodes is said to be "combinable" into a "macro-node" if the network remains acyclic after the combination of its members. For example, we cannot combine nodes A and H in the BN $\mathbb{A}$, as this results in the directed loop $\{A, H\} \rightarrow D \rightarrow \{A, H\}$. According to Chang & Fung's (1989) Node Aggregation Theorem, if all directed paths connecting any two nodes in a set do not contain a node outside it, then this set is combinable. Hence, set $\{A, C, D, H\}$ in the BN $\mathbb{A}$ is.

If two loops in a DAG share a common node, they belong to the same "independent loop set" (ILS). If we convert all ILSs in a DAG into polytrees, then the DAG itself becomes a polytree. Because each ILS is combinable, we can combine the nodes in each for this purpose. However this may yield some

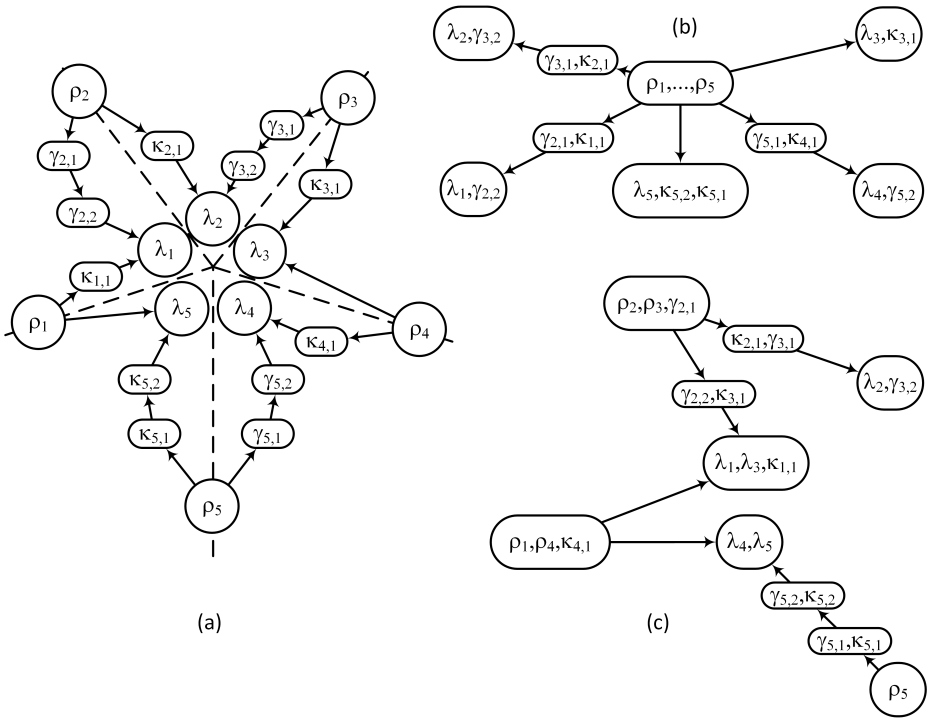

Figure 6: An isolated loop with five roots and five leaves

unnecessarily large macro-nodes. Chang & Fung suggest two methods of converting an ILS into a polytree; one is a heuristic search through the space of its "feasible partitions" for the optimal node combinations, the other is what they call the "graph-directed algorithm." Here, we propose an algorithm which is somewhat along the line of Ng and Levitt's method for incrementally extending what they called a "layered" polytree (1994).

6.1.1 The Isolated Loops

In a DAG, let us isolate one of its loops (that is, delete all nodes outside it) and call it an "isolated loop." Some non-root and non-leaf nodes of the DAG may now become the roots or leaves of the isolated loop. On the other hand, some roots or leaves of the former may no longer be in the latter.

Let us label the r roots in a loop consecutively as ρ_1, ..., ρ_r, with $\rho_{r+1} = \rho_1$. (See Figure 6(a).)

Lemma 19 *An isolated loop has the same number of roots and leaves, and can be organized into a star shape.*

Proof. Two neighboring roots cannot be in a parent-child relationship, thus must be connected from below via a leaf.

If there are $r = 2$ roots, then we need two leaves to have a loop connecting all four. However, if $r > 2$ and there are more than one leaves in the direct path connecting any two consecutive roots, we have a loop connecting all four, in addition to the loop connecting all r roots, which violates our isolated loop assumption. Thus with r roots, we have r leaves. Figure 6(a) shows a typical loop which is organized into a star shape. ■

Let λ_i be the single leaf between ρ_i and ρ_{i+1}. Let $\kappa_{i,k}$ $(k = 1, ..., n_i)$ be the n_i nodes in the directed path from ρ_i to λ_i and $\gamma_{i,k}$ $(k = 1, ..., m_i)$ be the m_i nodes in the directed path from ρ_{i+1} to λ_i.

There are more than one ways we can combine the nodes in an isolated loop in order to "open" it; that is, to make it a polytree. For example, we can combine all its roots after some appropriate node combinations to make path $\rho_i \to ... \to \lambda_i$ and path $\rho_{i+1} \to ... \to \lambda_i$ having the same "length;" that is, have the same number of macro-nodes. For example, to have the polytree in Figure 6(b), we form path $\rho_2 \to \gamma_{2,1} \to \{\lambda_1, \gamma_{2,2}\}$, so that it has the same length as path $\rho_1 \to \kappa_{1,1} \to \lambda_1$, before combining them into path $\{\rho_1, \rho_2\} \to \{\gamma_{2,1}, \kappa_{1,1}\} \to \{\lambda_1, \gamma_{2,2}\}$. Similarly, we can combine all its leaves.

Another novel method is to "fold" the loop along one of its axes, bringing together the pairs of opposite paths in order to combine them, creating an acyclic semi-star. The star must be made "symmetrical" around the chosen axis first, in that each pair of opposite paths must have the same length. The loop in Figure 6(a) is not symmetrical around axis $\lambda_2 - \rho_5$. To make it symmetrical, among other combinations, we may form path $\{\rho_2, \gamma_{2,1}\} \to \gamma_{2,2} \to \lambda_1$, so that it can be combined with path $\rho_3 \to \kappa_{3,1} \to \lambda_3$. The polytree in Figure 6(c) is one possible result of folding the loop in Figure 6(a) along its axis $\lambda_2 - \rho_5$.

So it seems that we can covert a DAG into a polytree by open all loops in isolation, one-by-one in this manner. Unfortunately, this procedure generally is not correct, as this may make other loops in the same ILS cyclic. For example, if there is an additional directed path $\rho_2 \to \eta \to \rho_3$ in Figure 6(a), then combining ρ_2 and ρ_3 as in Figure 6(c) violates the Node Aggregation Theorem. Furthermore, these cyclic loops may not be easily identifiable.

We now suggest a correct and systematic way to identify and open all loops in a DAG.

6.1.2 The Parentless Polytree Method

In our "parentless polytree method," we construct a growing parentless polytree from a DAG as follows: We start with a root in the DAG and execute the following steps:

1. While there is an un-recruited node
   (a) Recruit node τ in a topological order (that is, after all its parents $\{\pi_1, ..., \pi_p\}$), together with edge $\tau \leftarrow \pi_1$
   (b) While there is an un-recruited edge $\tau \leftarrow \pi_i$
      i. recruit edge $\tau \leftarrow \pi_i$
      ii. isolate and open any resulted loop
   (c) End while
   (d) Go to Step 1
2. End while

The growing polytree must be kept parentless, so that no directed path can return to it, and thus no directed cycle via a node outside the polytree can occur as a result of any node combinations inside the polytree.

When we recruit node τ having p parents into the polytree $\mathcal{P}$, we bring with it p edges $\tau \leftarrow \pi_i$ $(i = 1, ..., p)$, where each pair $\{\pi_i, \pi_j\}$ is connected by at most one path. (We have no path if π_i and π_j are in two unconnected subsets of $\mathcal{P}$.) Recruiting τ with all its parents creates at most C_2^p loops.

"Recruiting edge $\tau \leftarrow \pi_i$" means adding it in the polytree. Following Ng and Levitt (1994), not only do we recruit the nodes one-at-a-time, we also recruit the edges one-at-a-time. If τ has only one parent, then with edge $\tau \leftarrow \pi_1$ we still have a polytree. Otherwise, after recruiting edge $\tau \leftarrow \pi_2$, we have at most one loop $\tau \leftarrow \pi_1 - ... - \pi_2 \rightarrow \tau$. After opening this loop, we again have a polytree. Then edge $\tau \leftarrow \pi_3$ yields at most another loop via π_3.... In this way, we only have to open at most $p - 1$ easily identifiable loops.

If two loops created by τ have different pairs of parents $\{\pi_i, \pi_j\}$ and $\{\pi_k, \pi_m\}$, then opening one cannot make the other cyclic. So let us consider two loops having parents $\{\pi_i, \pi_j\}$ and $\{\pi_j, \pi_k\}$ and recruit edges $\tau \leftarrow \pi_i$ and $\tau \leftarrow \pi_j$ first, resulting in a loop having path $\pi_i \rightarrow \tau \leftarrow \pi_j$. (If $\mathcal{P}$ is not parentless, we may have a loop with path $\pi_i \rightarrow \tau \rightarrow \psi$ instead.) If we open this loop without combining τ with any node, then when we recruit edge $\tau \leftarrow \pi_k$ later, the resulting loop is acyclic, because path $\pi_j \rightarrow \tau \leftarrow \pi_k$ is still in it. It is feasible not to combine leaf τ in the folding method, as we do not need the macro-node $\{\tau, \pi_i\}$ to reduce the length of path $\tau \leftarrow \pi_i \leftarrow ... \leftarrow \rho_k$ and we can choose the symmetrical axis going through τ, so that it is not combined with any other leaf.

But suppose we wish to combine leaf τ with node η in the loop with $\{\pi_i, \pi_j\}$, and only to find out later when recruiting edge $\tau \leftarrow \pi_k$ that loop $\{\eta, \tau\} \rightarrow ... \rightarrow \pi_k \rightarrow \{\eta, \tau\}$ is cyclic because of the directed path $\eta \rightarrow ... \rightarrow \pi_k \rightarrow \tau$. In this case, we "open" this loop by expand the macro-node $\{\eta, \tau\}$ to $\{\eta, ..., \pi_k, \tau\}$. If there is any loop in this macro-node, it has the form $\tau \leftarrow ... - \eta \rightarrow ... \rightarrow \pi_k \rightarrow \tau$ and therefore is not cyclic. After recruiting all p edges, we have a polytree (without any cycle, hence is acyclic), and can return to Step 1.

After recruiting all nodes and their edges in a DAG in this manner, we will have partitioned the DAG into what we call a "macro-node polytree."

For illustration, let us use the BN in Figure 7, which we will refer to as the Bayesian network $\mathbb{C}$, or BN $\mathbb{C}$. Assume that, after forming the macro-nodes $\{R, T\}$ and $\{B, C\}$, we arrive at the parentless polytree above the curved line in Figure 7(a).

1. Recruiting node X results in loop $X - U - \{R, T\} - V - X$, which can be opened by forming the macro-node $\{U, V\}$.

2. Recruiting node Y and then node Z results in loop $Z - Y - \{U, V\} - X - Z$, which can be opened by forming the macro-node $\{X, Y\}$.

3. Recruiting node Q results in loop $Q - \{U, V\} - \{R, T\} - S - M - Q$, which can be opened by forming the macro-node $\{M, S, U, V\}$.

4. Recruiting edges $P \leftarrow \{M, S, U, P\}$ and $P \leftarrow Q$ results in loop $P - Q - \{M, S, U, V\} - P$, which can be opened by forming the macro-node $\{M, S, U, V, Q\}$.

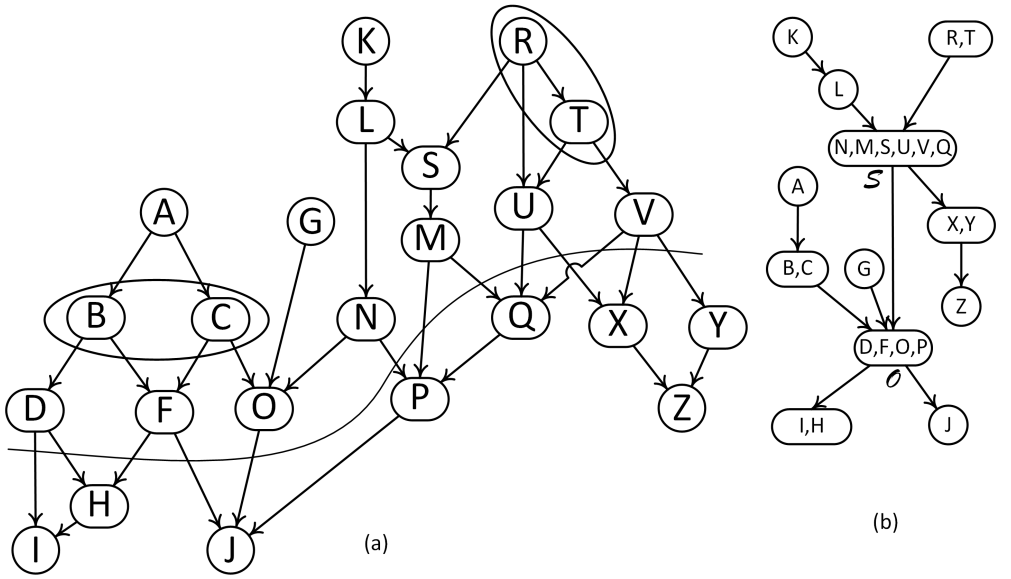

Figure 7: The Bayesian Network $\mathbb{C}$, or BN $\mathbb{C}$

5. Recruiting edge $P \leftarrow N$ results in loop $P-N-L-\{M,S,U,V,Q\}-P$, which can be opened by forming the macro-node $\{N,M,S,U,V,Q\}$.

6. Recruiting edges $J \leftarrow P$ and $J \leftarrow O$ results in loop $J-P-\{N,M,S,U,V,Q\}-O-J$, which can be opened by forming the macro-node $\{O,P\}$.

7. Recruiting edge $J \leftarrow F$ results in loop $J-\{O,P\}-\{B,C\}-F-J$, which can be opened by forming the macro-node $\{F,O,P\}$.

8. Recruiting node H results in loop $H-\{F,O,P\}-\{B,C\}-D-H$, which can be opened by forming the macro-node $\{D,F,O,P\}$.

9. Recruiting node I results in loop $I-H-\{D,F,O,P\}-I$, which can be opened by forming the macro-node $\{I,H\}$.

We now can apply the revised polytree algorithm to the final macro-node polytree corresponding to the BN $\mathbb{C}$, as shown in Figure 7(b). However, the size of some CPTs can be large, such as

$$\Pr\{D,F,O,P|N,M,S,U,V,Q,G,B,C\} = \Pr\{\mathcal{O}|\mathcal{S},G,B,C\}, \qquad (19)$$

where $\mathcal{O} = \{D, F, O, P\}$ and $\mathcal{S} = \{N, M, S, U, V, Q\}$. This is why we do not stop after Stage I, but continue to Stage II.

6.2 Stage II: The Border Polytrees

In Stage II, we use the border algorithm, which explores the independence relationships between the individual nodes within the macro-nodes, to stretch each macro-node into a border chain, if it is not already in this form.

A result obtained for the BN $\mathbb{C}$ after Stage II is shown in Figure 8.

Inside a macro-node, we follow the eight rules in §2.2 for promoting a node from a border and recruiting its cohort, so that each border separates the top part of the macro-node with its bottom part. However, because not all macro-notes are parentless, we need to introduce the following additional rules:

9. We cannot stretch a macro-node before its parents. In the BN $\mathbb{C}$, we must start with the macro-node $\{K\}$, $\{R, T\}$, $\{G\}$ or $\{A\}$.

10. Macro-node Γ cannot recruit any member of its macro-node child Δ. In the macro-node $\mathcal{S}$ of the BN $\mathbb{C}$, border $\{N, M, Q\}$ is obtained from border $\{N, M, Q, V\}$ by promoting V without recruiting its child Y in macro-node $\{X, Y\}$.

11. If macro-node Δ has parents in macro-node Γ, the entire set $\Gamma \cap \mathcal{H}_\Delta$ must be together in at least one border in Γ. In the BN $\mathbb{C}$, because $\mathcal{S} \cap \mathcal{H}_{\{X,Y\}} = \{U, V\}$ we form border $\{N, M, U, V\} \subseteq \mathcal{S}$, because $\mathcal{S} \cap \mathcal{H}_{\mathcal{O}} = \{N, M, Q\}$ we form border $\{N, M, Q\} \subseteq \mathcal{S}$. Also, because $\mathcal{O} \cap \mathcal{H}_{\{H,I\}} = \{D, F\}$ and $\mathcal{O} \cap \mathcal{H}_J = \{F, O, P\}$, we form border $\{D, F, O, P\} \subseteq \mathcal{O}$.

12. Suppose we wish to promote node $B \in \mathcal{B}_{i-1} \subseteq \Delta$. By Rule 10, we do not recruit $\mathcal{L}_B \nsubseteq \Delta$. However, we need to recruit all un-recruited members of $\mathcal{L}_B \cap \Delta$ (that is, $\mathcal{F}_B = \mathcal{L}_B \cap \mathcal{D}_{\mathcal{B}_{i-1}} \cap \Delta$) and all members of $\mathcal{H}_{\mathcal{F}_B} \subseteq \mathcal{V}$ (that is, the co-parents of B), even if they are not in Δ.

    Assume there are r macro-nodes $\Gamma_k \neq \Delta$ $(k = 1, ..., r)$ such that $\Gamma_k \cap \mathcal{H}_{\mathcal{F}_B} \neq \varnothing$. By Rules 9 and 11, we must have already constructed in Γ_k a border $\mathcal{B}_{i-1,k}$ such that $\{\Gamma_k \cap \mathcal{H}_{\mathcal{F}_B}\} \subseteq \mathcal{B}^*_{i-1,k} = \{\Gamma_k \cap \mathcal{H}_\Delta\} \subseteq \mathcal{B}_{i-1,k}$. For notational advantages, we also denote $\mathcal{B}_{i-1} \subseteq \Delta$ by $\mathcal{B}_{i-1,0}$ or $\mathcal{B}^*_{i-1,0}$.

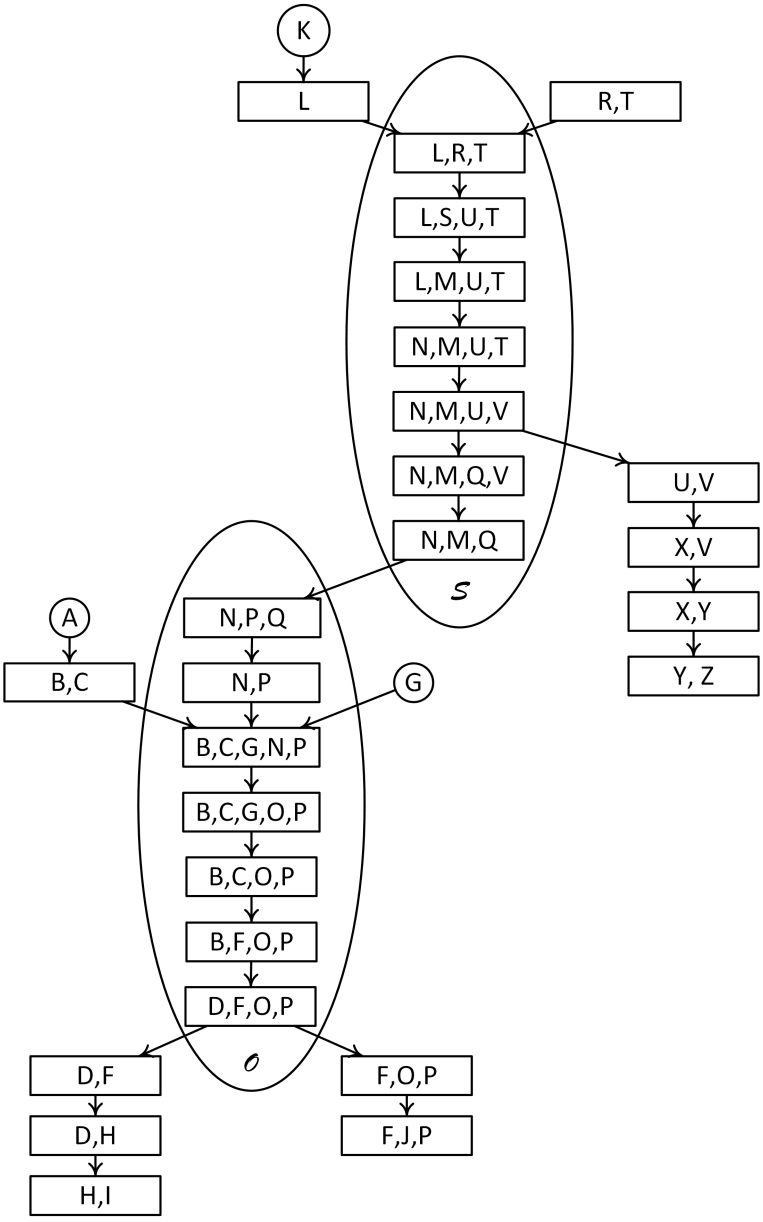

Figure 8: BP $\mathbb{C}$, a border polytree for the BN $\mathbb{C}$

We break our procedure into two steps: (i) We first recruit to Δ all $\Gamma_k \cap \mathcal{H}_\Delta$ $(k = 1, ..., r)$ (which have not been recruited) to form border

$$\mathcal{B}_i = \mathcal{B}_{i-1} \cup_{j=1}^{r} \{\Gamma_k \cap \mathcal{H}_\Delta\} = \cup_{j=0}^{r} \mathcal{B}^*_{i-1,j} \supseteq \mathcal{H}_{\mathcal{F}_B}, \tag{20}$$

where $\mathcal{B}_i$ has all $\mathcal{B}_{i-1,k}$ as parents. (ii) Then we promote node B from border $\mathcal{B}_i$ with cohort $\mathcal{F}_B \subseteq \Delta$ by Rule 2.

Note that the entire set $\Gamma_k \cap \mathcal{H}_\Delta$ constructed by Rule 11 must be recruited into Δ together. This is to ensure that there is at most one edge connecting any two stretched macro-nodes. Otherwise, the main benefit of the macro-node polytrees is destroyed.

In the BN $\mathbb{C}$, suppose we wish to promote node N from border $\{N, P\} \subset \mathcal{O}$. Because $\mathcal{F}_N = \mathcal{L}_N \cap \mathcal{D}_{\{N,P\}} \cap \mathcal{O} = O$, we need to recruit nodes $C \in \{B, C\} \cap \mathcal{H}_O$ and $G \in \{G\} \cap \mathcal{H}_O$. However, Rule 11 puts $\{B, C\} \cap \mathcal{H}_\mathcal{O} = \{B, C\}$ in one border and instead of C alone, we recruit the whole set $\{B, C\}$. We thus form border $\cup_{j=0}^{2} \mathcal{B}^*_{i-1,j}$ with $\mathcal{B}^*_{i-1,0} = \{N, P\}$, $\mathcal{B}^*_{i-1,1} = \{B, C\}$, $\mathcal{B}^*_{i-1,2} = \{G\}$ before promoting N with cohort O.

13. Rule 12 also applies when we start stretching a non-root macro-node Δ. We do not start with the nodes having parents in it, but with node V having parents only in other macro-nodes $\Gamma_k \neq \Delta$ $(k = 1, ..., r)$; that is, $\Delta \cap \mathcal{H}_V = \varnothing$, and $\Gamma_k \cap \mathcal{H}_V \neq \varnothing$. In this case we first form border $\mathcal{B}_i$ as in Equation (20), with $\mathcal{B}_{i-1,0} = \mathcal{B}^*_{i-1,0} = \varnothing$ and with $\mathcal{B}^*_{i-1,k}$ such that $\{\Gamma_k \cap \mathcal{H}_V\} \subseteq \mathcal{B}^*_{i-1,k} = \{\Gamma_k \cap \mathcal{H}_\Delta\} \subseteq \mathcal{B}_{i-1,k}$. Then we can recruit V.

    In the BN $\mathbb{C}$, to start stretching macro-node $\mathcal{S}$ with node S having parents in macro-nodes $\{L\}$ and $\{R, T\}$, we first form border $\cup_{j=1}^{2} \mathcal{B}^*_{i-1,j}$ with $\mathcal{B}^*_{i-1,1} = \mathcal{H}_\mathcal{S} \cap \{L\} = L$ and $\mathcal{B}^*_{i-1,2} = \mathcal{H}_\mathcal{S} \cap \{R, T\} = \{R, T\}$. We then can promote R with cohort $\{S, U\}$.

    If $r = 1$ and $\mathcal{B}^*_{i-1,1} = \mathcal{B}_{i-1,1}$, then there is no need to repeat $\mathcal{B}_i = \mathcal{B}_{i-1,1}$. Suppose we wish to start stretching macro-node $\mathcal{O}$ in the BN $\mathbb{C}$ by recruiting node P. Because P only has parents in $\mathcal{S}$, there is no need to repeat border $\mathcal{B}^*_{i-1,1} = \mathcal{B}_{i-1,1} = \mathcal{S} \cap \mathcal{H}_\mathcal{O} = \{N, M, Q\}$. We simply promote M to recruit P.

After stretching all macro-nodes, we obtain a "border polytree" or a BP, made up of the "borders." The variables in each border are called its "component variables." The BP in Figure 8 is called the BP $\mathbb{C}$.

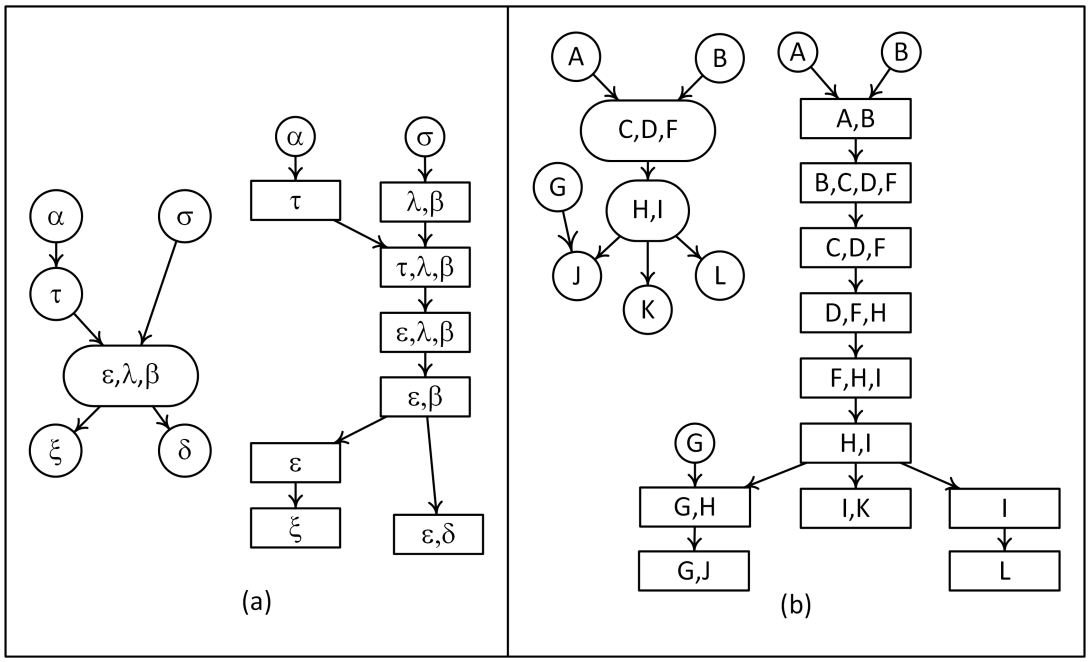

Figure 9: The macro-node & border polytrees for (a) the Dyspnoea example and (b) the BN $\mathbb{A}$

For comparison with a undirected junction tree of the well-known "Dyspnoea" example by Lauritzen & Spiegelhalter (1988, pp. 164 and 209), we include in Figure 9(a) its macro-node and border polytrees. Figure 9(b) shows the macro-node and border polytrees of the BN $\mathbb{A}$.

Like a clique tree (or a junction tree), a BP has the "running intersection property;" that is, after a variable is recruited into border $\mathcal{B}_j$, it stays in all future consecutive borders $\mathcal{B}_{j+1}$, $\mathcal{B}_{j+2}$,... until its promotion; then it will never be recruited again into any later borders. However, a BP is not "family preserving," in that it is not necessary that a variable must be with all its parents in one border. Most importantly, a BP is directional, while a clique tree is not.

The borders in a BP generally have smaller CPTs than the macro-nodes in macro-node polytree. In the BN $\mathbb{C}$, instead of the CPT as in Equation (19) involving thirteen variables, we now have $\Pr\{D, F, O, P|B, F, O, P\}$, with five variables $\{D, F, O, P, B\}$. Despite their smaller CPTs, their poste-

rior marginal probabilities still can be calculated by adapting the results in Section 4 as we will show next.

6.3 The Message Propagations in a Border Polytree

There are two types of non-root borders in a BP:

1. $\mathcal{B}_i = \{\mathcal{B}_{i-1}, \mathcal{C}_i\} \setminus V_i$, obtained by following Rules 1-8 in §2.2, having $\mathcal{B}_{i-1}$ as a single parent. (If $\mathcal{B}_i = X$ and $\mathcal{B}_{i-1} = Y$, then $\mathcal{C}_i = X$, $V_i = Y$.)
2. $\mathcal{B}_i = \cup_{j=0}^{r} \mathcal{B}^*_{i-1,j}$, obtained by Rule 12 or 13, with $\mathcal{H}_{\mathcal{B}_i} = \{\mathcal{B}_{i-1,j}, j = 0, ..., r\}$.

Returning to the proofs of Lemmas 10 and 13, we see that they are still valid for the BPs. However, we need to modify Lemmas 11 and 12 according to the type of border in a BP.

6.3.1 The Downward Propagations

We first modify Lemma 12:

Lemma 20 *For border* $\mathcal{B}_i = \{\mathcal{B}_{i-1}, \mathcal{C}_i\} \setminus V_i$,

$$\Pi(\mathcal{B}_i) = \sum_{V_i} \phi(\mathcal{C}_i) \Pi_{\mathcal{B}_i}(\mathcal{B}_{i-1}).$$

If $\mathcal{B}_i$ *is the only child of* $\mathcal{B}_{i-1}$,

$$\Pi(\mathcal{B}_i) = \sum_{V_i} \phi(\mathcal{C}_i) \Pi(\mathcal{B}_{i-1}).$$

Proof. Consider the parentless $\mathcal{T}_{\mathcal{B}_{i-1} \to \mathcal{B}_i}$ with border $\mathcal{B}_{i-1}$. Promoting $V_i \in \mathcal{B}_{i-1}$ with cohort $\mathcal{C}_i$ yields the parentless $\mathcal{P}_{\mathcal{B}_i}$ with border $\mathcal{B}_i$. This because, with $\{\mathcal{B}_{i-1} \setminus V_i\} \subseteq \mathcal{T}_{\mathcal{B}_{i-1} \to \mathcal{B}_i}$,

$$\mathcal{T}_{\mathcal{B}_{i-1} \to \mathcal{B}_i} \cup \mathcal{C}_i = \mathcal{T}_{\mathcal{B}_{i-1} \to \mathcal{B}_i} \cup \{\mathcal{B}_{i-1} \setminus V_i\} \cup \mathcal{C}_i = \mathcal{T}_{\mathcal{B}_{i-1} \to \mathcal{B}_i} \cup \mathcal{B}_i = \mathcal{P}_{\mathcal{B}_i}.$$

The lemma follows from Theorem 5.

If $\mathcal{B}_i$ is the only child of $\mathcal{B}_{i-1}$, by Lemma 10, $\Pi_{\mathcal{B}_i}(\mathcal{B}_{i-1}) = \Pi(\mathcal{B}_{i-1})$. ■

Lemma 21 *For border* $\mathcal{B}_i = \cup_{j=0}^{r} \mathcal{B}_{i-1,j}^{*}$,

$$\Pi\left(\mathcal{B}_i\right) = \prod_{j=0}^{r} \sum_{\mathcal{B}_{i-1,j} \backslash \mathcal{B}_{i-1,j}^{*}} \Pi_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,j}\right).$$

Proof. Consider the parentless $\overline{\mathcal{T}}_{i-1} = \cup_{j=0}^{r} \mathcal{T}_{\mathcal{B}_{i-1,j} \rightarrow \mathcal{B}_i}$ having border $\overline{\mathcal{B}}_{i-1} = \cup_{j=0}^{r} \mathcal{B}_{i-1,j}$. Because all $\mathcal{T}_{\mathcal{B}_{i-1,j} \rightarrow \mathcal{B}_i}$ are parentless and independent, from Definitions (9) and (15),

$$\begin{aligned}
\Pi\left(\overline{\mathcal{B}}_{i-1}\right) &= \Pr\left\{\overline{\mathcal{B}}_{i-1}, \left[\overline{\mathcal{T}}_{i-1} \backslash \overline{\mathcal{B}}_{i-1}\right]\right\} \boldsymbol{I}_{\overline{\mathcal{B}}_{i-1}} \\
&= \Pr\left\{\cup_{j=0}^{r}\left\{\mathcal{B}_{i-1,j}, \left[\mathcal{T}_{\mathcal{B}_{i-1,j} \rightarrow \mathcal{B}_i} \backslash \mathcal{B}_{i-1,j}\right]\right\}\right\} \boldsymbol{I}_{\cup_{j=0}^{r} \mathcal{B}_{i-1,j}} \\
&= \prod_{j=0}^{r} \Pr\left\{\mathcal{B}_{i-1,j}, \left[\mathcal{T}_{\mathcal{B}_{i-1,j} \rightarrow \mathcal{B}_i} \backslash \mathcal{B}_{i-1,j}\right]\right\} \boldsymbol{I}_{\mathcal{B}_{i-1,j}} = \prod_{j=0}^{r} \Pi_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,j}\right).
\end{aligned}$$

Because $\mathcal{B}_{i-1,j} \backslash \mathcal{B}_{i-1,j}^{*}$ does not have children in the bottom set $\{\mathcal{B}_i, \mathcal{D}_{\mathcal{B}_i}\}$, we use Rule 1 to promote $\cup_{j=0}^{r}\left\{\mathcal{B}_{i-1,j} \backslash \mathcal{B}_{i-1,j}^{*}\right\}$ without cohort to obtain the same parentless $\overline{\mathcal{T}}_{i-1}$ with border $\mathcal{B}_i = \cup_{j=0}^{r} \mathcal{B}_{i-1,j}^{*}$. From Theorem 5,

$$\Pi\left(\mathcal{B}_i\right) = \sum_{\cup_{j=0}^{r}\left\{\mathcal{B}_{i-1,j} \backslash \mathcal{B}_{i-1,j}^{*}\right\}} \Pi\left(\overline{\mathcal{B}}_{i-1}\right) = \sum_{\cup_{j=0}^{r}\left\{\mathcal{B}_{i-1,j} \backslash \mathcal{B}_{i-1,j}^{*}\right\}} \prod_{j=0}^{r} \Pi_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,j}\right),$$

hence the lemma. ■

Combining Lemma 10 with above lemmas yields the following BP version of Theorem 14.

Theorem 22 *If border* $\mathcal{B}_i$ *has received the messages from all members of* $\mathcal{H}_{\mathcal{B}_i} \cup \{\mathcal{L}_{\mathcal{B}_i} \backslash \mathcal{B}_{i+1,j}\}$, *then it can send a downward message to its child* $\mathcal{B}_{i+1,j}$ *as*

$$\Pi_{\mathcal{B}_{i+1,j}}\left(\mathcal{B}_i\right) = \Pi\left(\mathcal{B}_i\right) \left(\prod_{\mathcal{W} \in \mathcal{L}_{\mathcal{B}_i} \backslash \mathcal{B}_{i+1,j}} \Lambda_{\mathcal{W}}\left(\mathcal{B}_i\right) \right),$$

where $\Pi\left(\mathcal{B}_i\right)$ *is calculated by Lemma 20 or 21.*

6.3.2 The Upward Propagations

We now present the following BP versions of Lemma 11:

Lemma 23 *For border* $\mathcal{B}_i = \{\mathcal{B}_{i-1}, \mathcal{C}_i\} \setminus V_i$,

$$\Lambda_{\mathcal{B}_i}(\mathcal{B}_{i-1}) = \sum_{\mathcal{C}_i} \phi(\mathcal{C}_i) \Lambda(\mathcal{B}_i).$$

If $\mathcal{B}_i$ *is the only child of* $\mathcal{B}_{i-1}$,

$$\Lambda(\mathcal{B}_{i-1}) = \sum_{\mathcal{C}_i} \phi(\mathcal{C}_i) \Lambda(\mathcal{B}_i).$$

Proof. Consider the parentless $\mathcal{T}_{\mathcal{B}_{i-1} \to \mathcal{B}_i}$ with border $\mathcal{B}_{i-1}$. As shown in Lemma 20, promoting $V_i \in \mathcal{B}_{i-1}$ with cohort $\mathcal{C}_i$ yields the parentless $\mathcal{P}_{\mathcal{B}_i}$ with border $\mathcal{B}_i$. The lemma follows from Theorem 7.

If $\mathcal{B}_i$ is the only child of $\mathcal{B}_{i-1}$, by Lemma 13, $\Lambda_{\mathcal{B}_i}(\mathcal{B}_{i-1}) = \Lambda(\mathcal{B}_{i-1})$. ■

Lemma 24 *If border* $\mathcal{B}_i = \cup_{j=0}^{r} \mathcal{B}_{i-1,j}^*$, *then for all* $0 \le k \le r$,

$$\Lambda_{\mathcal{B}_i}(\mathcal{B}_{i-1,k}) = \sum_{\mathcal{H}_{\mathcal{B}_i} \setminus \mathcal{B}_{i-1,k}} \Lambda(\mathcal{B}_i) \prod_{\mathcal{W} \in \mathcal{H}_{\mathcal{B}_i} \setminus \mathcal{B}_{i-1,k}} \Pi_{\mathcal{B}_i}(\mathcal{W}).$$

Proof. Consider the parentless $\mathcal{T}_{\mathcal{B}_{i-1,k} \to \mathcal{B}_i}$ $(0 \le k \le r)$ having border $\mathcal{B}_{i-1,k}$. Let $\mathcal{I}_{i-1,k} = \mathcal{H}_{\mathcal{B}_i} \setminus \mathcal{B}_{i-1,k}$. We recruit $\cup_{\mathcal{W} \in \mathcal{I}_{i-1,k}} \mathcal{T}_{\mathcal{W} \to \mathcal{B}_i}$ without promotion, resulting in the parentless $\overline{\mathcal{T}}_{i-1} = \cup_{j=0}^{r} \mathcal{T}_{\mathcal{B}_{i-1,j} \to \mathcal{B}_i}$ having border $\overline{\mathcal{B}}_{i-1} = \cup_{j=0}^{r} \mathcal{B}_{i-1,j}$. From Theorem 7,

$$\begin{aligned}
&\Lambda_{\mathcal{B}_i}(\mathcal{B}_{i-1,k}) \\
&= \sum_{\cup_{\mathcal{W} \in \mathcal{I}_{i-1,k}} \mathcal{T}_{\mathcal{W} \to \mathcal{B}_i}} \left(\prod_{\mathcal{W} \in \mathcal{I}_{i-1,k}} \Pr\{\mathcal{T}_{\mathcal{W} \to \mathcal{B}_i}\} \boldsymbol{I}_{\mathcal{T}_{\mathcal{W} \to \mathcal{B}_i}} \right) \Lambda\left(\overline{\mathcal{B}}_{i-1}\right) \\
&= \sum_{\mathcal{I}_{i-1,k}} \left\{ \sum_{\cup_{\mathcal{W} \in \mathcal{I}_{i-1,k}} \{\mathcal{T}_{\mathcal{W} \to \mathcal{B}_i} \setminus \mathcal{W}\}} \left(\prod_{\mathcal{W} \in \mathcal{I}_{i-1,k}} \Pr\{\mathcal{T}_{\mathcal{W} \to \mathcal{B}_i}\} \boldsymbol{I}_{\mathcal{T}_{\mathcal{W} \to \mathcal{B}_i}} \right) \Lambda\left(\overline{\mathcal{B}}_{i-1}\right) \right\}.
\end{aligned}$$

Because $\overline{\mathcal{B}}_{i-1} \cap \left\{\cup_{\mathcal{W}\in\mathcal{I}_{i-1,k}} \mathcal{T}_{\mathcal{W}\to\mathcal{B}_i} \backslash \mathcal{W}\right\} = \varnothing$,

$$\begin{aligned}
&\Lambda_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,k}\right) \\
&= \sum_{\mathcal{I}_{i-1,k}} \Lambda\left(\overline{\mathcal{B}}_{i-1}\right) \left\{ \sum_{\cup_{\mathcal{W}\in\mathcal{I}_{i-1,k}} \left\{\mathcal{T}_{\mathcal{W}\to\mathcal{B}_i}\backslash\mathcal{W}\right\}} \left(\prod_{\mathcal{W}\in\mathcal{I}_{i-1,k}} \Pr\left\{\mathcal{T}_{\mathcal{W}\to\mathcal{B}_i}\right\} \boldsymbol{I}_{\mathcal{T}_{\mathcal{W}\to\mathcal{B}_i}} \right) \right\} \\
&= \sum_{\mathcal{I}_{i-1,k}} \Lambda\left(\overline{\mathcal{B}}_{i-1}\right) \prod_{\mathcal{W}\in\mathcal{I}_{i-1,k}} \left(\sum_{\mathcal{T}_{\mathcal{W}\to\mathcal{B}_i}\backslash\mathcal{W}} \Pr\left\{\mathcal{T}_{\mathcal{W}\to\mathcal{B}_i}\right\} \boldsymbol{I}_{\mathcal{T}_{\mathcal{W}\to\mathcal{B}_i}} \right).
\end{aligned}$$

By Definition (15),

$$\Lambda_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,k}\right) = \sum_{\mathcal{I}_{i-1,k}} \Lambda\left(\overline{\mathcal{B}}_{i-1}\right) \prod_{\mathcal{W}\in\mathcal{I}_{i-1,k}} \Pi_{\mathcal{B}_i}\left(\mathcal{W}\right).$$

Now use Rule 1 to promote $\cup_{j=0}^{r}\left\{\mathcal{B}_{i-1,j}\backslash\mathcal{B}_{i-1,j}^{*}\right\}$ without cohort, leaving border $\mathcal{B}_i = \cup_{j=0}^{r}\mathcal{B}_{i-1,j}^{*}$. From Theorem 7, $\Lambda\left(\overline{\mathcal{B}}_{i-1}\right) = \Lambda\left(\mathcal{B}_i\right)$, hence the lemma. ■

Lemma 24 can be carried out more efficiently as

$$\begin{aligned}
&\Lambda_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,k}\right) \\
&= \sum_{\mathcal{B}_{i-1,r}} \Pi_{\mathcal{B}_i}^{\mathcal{B}_{i-1,r}} ... \sum_{\mathcal{B}_{i-1,k+1}} \Pi_{\mathcal{B}_i}^{\mathcal{B}_{i-1,k+1}} \sum_{\mathcal{B}_{i-1,k-1}} \Pi_{\mathcal{B}_i}^{\mathcal{B}_{i-1,k-1}} ... \sum_{\mathcal{B}_{i-1,0}} \Pi_{\mathcal{B}_i}^{\mathcal{B}_{i-1,0}} \Lambda\left(\mathcal{B}_i\right),
\end{aligned}$$

where $\Pi_{\mathcal{B}_i}^{\mathcal{B}_{i-1,j}} = \Pi_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,j}\right)$ for all $0 \leq j \leq r$.

Combining Lemma 13 with the above lemmas yields the BP version of Theorem 15:

Theorem 25 *If border $\mathcal{B}_i$ has received the messages from all members of $\{\mathcal{H}_{\mathcal{B}_i}\backslash\mathcal{B}_{i-1,k}\} \cup \mathcal{L}_{\mathcal{B}_i}$, then it can send an upward message $\Lambda_{\mathcal{B}_i}\left(\mathcal{B}_{i-1,k}\right)$ to its parent $\mathcal{B}_{i-1,k}$ as in Lemma 23 or 24, in which*

$$\Lambda\left(\mathcal{B}_i\right) = \prod_{\mathcal{W}\in\mathcal{L}_{\mathcal{B}_i}} \Lambda_{\mathcal{W}}\left(\mathcal{B}_i\right).$$

Both the downward and upward passes use the reduced cohort probability tables $\phi\left(\cdot\right)$, rather than the individual CPTs. So all $\Phi\left(\cdot\right)$ should be calculated and pre-loaded.

6.3.3 The Border Evidential Cores

A border is evidential if one of its component variables is. Let the "border evidential core" (or "border EC") be the *smallest* sub-polytree which contains all the evidence variables in its borders. For the BN $\mathbb{C}$, let us assume $\mathcal{E} = \{B, O, Q\}$. Then the largest sub-polytree in Figure 8 containing all the evidence variables includes path $\{N, M, Q, V\} \rightarrow \{N, M, Q\} \rightarrow \{N, P, Q\} \rightarrow ... \rightarrow \{D, F, O, P\}$ and path $\{B, C\} \rightarrow \{B, C, G, N, P\}$. However, its border EC is smaller, including only path $\{N, P, Q\} \rightarrow \{N, P\} \rightarrow \{B, C, G, N, P\} \rightarrow \{B, C, G, O, P\}$. For some evidence sets $\mathcal{E}$, the border EC is not unique.

A border is "inside" if it is in the border EC; otherwise it is "outside;" edge $\mathcal{B}_i - \mathcal{B}_j$ is "inside" if both borders are inside; otherwise it is "outside;" a path is "inside" if all its component edges are.

6.3.4 The Boundary Conditions

As in the polytrees, the calculations of the messages along an inside edge of a border EC may require the knowledge of the boundary conditions; that is, the messages to an inside border from its outside neighboring borders. However, unlike the polytree, our definition of the border ECs allows the outside borders to be evidential. (For example, borders $\{N, M, Q, V\}$ or $\{B, F, O, P\}$ in Figure 8 with $\mathcal{E} = \{B, O, Q\}$.) So we need the following BP version of Theorem 16:

Theorem 26 *Consider an inside border $\mathcal{B}_i$.*
(a) If $\mathcal{B}_{i-1} \in \mathcal{H}_{\mathcal{B}_i}$ is outside, then $\Pi_{\mathcal{B}_i}(\mathcal{B}_{i-1}) = \Pr\{\mathcal{B}_{i-1}\} \boldsymbol{I}_{\mathcal{B}_{i-1}}$.
(b) If $\mathcal{B}_{i+1} \in \mathcal{L}_{\mathcal{B}_i}$ is outside, then $\Lambda_{\mathcal{B}_{i+1}}(\mathcal{B}_i) = \boldsymbol{I}_{\mathcal{B}_i}$.

Proof. (a) In a border polytree, assume $\mathcal{B}_i$ is inside, $\mathcal{B}_{i-1} \in \mathcal{H}_{\mathcal{B}_i}$ is outside, and the outside parentless $\mathcal{T}_{\mathcal{B}_{i-1} \rightarrow \mathcal{B}_i}$ has an evidence variable $E \in \widetilde{\mathcal{B}}_1$ such that $E \notin \mathcal{B}_i$. By its definition, the border EC must include another evidence border $\widetilde{\mathcal{B}}_2 \ni E$. By the running intersection property, the path connecting $\widetilde{\mathcal{B}}_1$ and $\widetilde{\mathcal{B}}_2$ cannot go through $\mathcal{B}_i$. Thus there are two paths connecting $\widetilde{\mathcal{B}}_1$ and $\widetilde{\mathcal{B}}_2$, including the one via $\mathcal{B}_i$. Because this is contradictory to the polytree assumption, we must have $E \in \mathcal{B}_i$. By the running intersection property, $E \in \mathcal{B}_{i-1}$. In other words, $\boldsymbol{I}_{\mathcal{T}_{\mathcal{B}_{i-1} \rightarrow \mathcal{B}_i}} = \boldsymbol{I}_{\mathcal{B}_{i-1}}$. By Definition (15),

$$\Pi_{\mathcal{B}_i}(\mathcal{B}_{i-1}) = \sum_{\mathcal{T}_{\mathcal{B}_{i-1} \rightarrow \mathcal{B}_i} \backslash \mathcal{B}_{i-1}} \Pr\left\{\mathcal{T}_{\mathcal{B}_{i-1} \rightarrow \mathcal{B}_i}\right\} \boldsymbol{I}_{\mathcal{T}_{\mathcal{B}_{i-1} \rightarrow \mathcal{B}_i}} = \Pr\{\mathcal{B}_{i-1}\} \boldsymbol{I}_{\mathcal{B}_{i-1}}.$$

(b) In a border polytree, assume $\mathcal{B}_i$ is inside, $\mathcal{B}_{i+1} \in \mathcal{L}_{\mathcal{B}_i}$ is outside, and the bottom part $\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}}$ has an evidence variable $E \in \widetilde{\mathcal{B}}_3$ such that $E \notin \mathcal{B}_i$. The border EC must include another evidence border $\widetilde{\mathcal{B}}_4 \ni E$. By the running intersection property, the path connecting $\widetilde{\mathcal{B}}_3$ and $\widetilde{\mathcal{B}}_4$ cannot go through $\mathcal{B}_i$. Thus there are two paths connecting $\widetilde{\mathcal{B}}_3$ and $\widetilde{\mathcal{B}}_4$, including the one via $\mathcal{B}_i$. As this is a contradiction to the polytree assumption, all evidence variables in $\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}}$ must be in $\mathcal{B}_i$. In other words, $\boldsymbol{I}_{\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}}} = \boldsymbol{I}_{\mathcal{B}_i}$. By Definition (18),

$$\begin{aligned}\Lambda_{\mathcal{B}_{i+1}}(\mathcal{B}_i) &= \sum_{\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}}} \Pr\left\{\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}} | \mathcal{B}_i\right\} \boldsymbol{I}_{\mathcal{B}_i \cup \mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}}} \\ &= \sum_{\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}}} \Pr\left\{\mathcal{U}_{\mathcal{B}_i \to \mathcal{B}_{i+1}} | \mathcal{B}_i\right\} \boldsymbol{I}_{\mathcal{B}_i} = \boldsymbol{I}_{\mathcal{B}_i}.\end{aligned}$$

■

6.3.5 The Message Initializations

Following is the BP version of Theorem 17:

Theorem 27 *Consider a border EC.*
(a) If $\mathcal{B}_R$ is a root with inside child $\mathcal{B}_C$, then $\Pi(\mathcal{B}_R) = \Pi_{\mathcal{B}_C}(\mathcal{B}_R) = \Pr\{\mathcal{B}_R\} \boldsymbol{I}_{\mathcal{B}_R}$.
(b) If $\mathcal{B}_L$ is a leaf, then $\Lambda(\mathcal{B}_L) = \boldsymbol{I}_{\mathcal{B}_L}$ in Lemma 23 or 24.

Proof. (a) Let $\mathcal{B}_R$ be a root of the border EC. As shown in the proof of Theorem 26(a), for all $\mathcal{W} \in \mathcal{H}_{\mathcal{B}_R}$, all evidence variables in $\mathcal{T}_{\mathcal{W} \to \mathcal{B}_R}$ are in the inside $\mathcal{B}_R$; so are all evidence variables in $\mathcal{P}_{\mathcal{B}_R} = \mathcal{B}_R \cup_{\mathcal{W} \in \mathcal{H}_{\mathcal{B}_R}} \mathcal{T}_{\mathcal{W} \to \mathcal{B}_R}$. Hence $\boldsymbol{I}_{\mathcal{P}_{\mathcal{B}_R}} = \boldsymbol{I}_{\mathcal{B}_R}$. From Definition (13),

$$\Pi(\mathcal{B}_R) = \sum_{\mathcal{A}_{\mathcal{B}_R}} \Pr\{\mathcal{P}_{\mathcal{B}_R}\} \boldsymbol{I}_{\mathcal{P}_{\mathcal{B}_R}} = \sum_{\mathcal{A}_{\mathcal{B}_R}} \Pr\{\mathcal{A}_{\mathcal{B}_R}, \mathcal{B}_R\} \boldsymbol{I}_{\mathcal{B}_R} = \Pr\{\mathcal{B}_R\} \boldsymbol{I}_{\mathcal{B}_R}.$$

Also, because $\mathcal{B}_R$ has only one inside child $\mathcal{B}_C$, $\Lambda_{\mathcal{W}}(\mathcal{B}_R) = \boldsymbol{I}_{\mathcal{B}_R}$ for all $\mathcal{W} \in \mathcal{L}_{\mathcal{B}_R} \backslash \mathcal{B}_C$. Thus $\Pi(\mathcal{B}_R) = \Pi_{\mathcal{B}_C}(\mathcal{B}_R)$ from Lemma 10.

(b) Let $\mathcal{B}_L$ be a leaf of the border EC. As shown in the proof of Theorem 26(b), for all $\mathcal{W} \in \mathcal{L}_{\mathcal{B}_L}$, all evidence variables in $\mathcal{U}_{\mathcal{B}_L \to \mathcal{W}}$ must be in $\mathcal{B}_L$; so are all evidence variables in $\mathcal{D}_{\mathcal{B}_L} = \cup_{\mathcal{W} \in \mathcal{L}_{\mathcal{B}_L}} \mathcal{U}_{\mathcal{B}_L \to \mathcal{W}}$. Hence $\boldsymbol{I}_{\mathcal{D}_{\mathcal{B}_L}} = \boldsymbol{I}_{\mathcal{B}_L}$. From Definition (14),

$$\Lambda(\mathcal{B}_L) = \sum_{\mathcal{D}_{\mathcal{B}_L}} \Pr\{\mathcal{D}_{\mathcal{B}_L} | \mathcal{B}_L\} \boldsymbol{I}_{\mathcal{B}_L \cup \mathcal{D}_{\mathcal{B}_L}} = \sum_{\mathcal{D}_{\mathcal{B}_L}} \Pr\{\mathcal{D}_{\mathcal{B}_L} | \mathcal{B}_L\} \boldsymbol{I}_{\mathcal{B}_L} = \boldsymbol{I}_{\mathcal{B}_L}.$$

■

For the border chain obtained in Section 2, recall that we denote the first time and the last time an evidence is recruited into $\mathcal{P}$ by α and β, respectively. We now see that the part of the border chain from α to β is its border EC, and Theorem 27 is consistent with Equations (10) and (12).

6.3.6 The Collection Phase

The inferences in a border polytree are carried out in the same manner as in a polytree: In the collection phase, we start with Theorem 27 and propagate the messages inside the border EC to an arbitrarily chosen inside pivot border.

For the border EC in Figure 8 with $\mathcal{E} = \{B, O, Q\}$, suppose we pick its leaf $\{B, C, G, O, P\}$ as the pivot, then the collection phase includes:

1. By Theorem 27(a),

$$\Pi(N, P, Q) = \Pr\{N, P, Q\} \boldsymbol{I}_Q.$$

2. By Lemma 20(b) and Theorem 22,

$$\Pi_{\{B,C,G,N,P\}}(N, P) = \Pi(N, P) = \sum_Q \Pr\{N, P, Q\} \boldsymbol{I}_Q.$$

3. By Lemma 21, with the boundary conditions $\Pi_{\{B,C,G,N,P\}}(B, C) = \Pr\{B, C\} \boldsymbol{I}_B$ and $\Pi_{\{B,C,G,N,P\}}(G) = \Pr\{G\}$,

$$\begin{aligned}
&\Pi(B, C, G, N, P) \\
&= \Pi_{\{B,C,G,N,P\}}(B, C)\, \Pi_{\{B,C,G,N,P\}}(G)\, \Pi_{\{B,C,G,N,P\}}(N, P) \\
&= \Pr\{B, C\} \boldsymbol{I}_B \Pr\{G\} \sum_Q \Pr\{N, P, Q\} \boldsymbol{I}_Q.
\end{aligned}$$

4. By Lemma 20(b), at the pivot border,

$$\begin{aligned}
\Pi(B, C, G, O, P) &= \sum_N \Pr\{O|C, G, N\} \boldsymbol{I}_O \Pi(B, C, G, N, P) \\
&= \sum_N \Pr\{O|C, G, N\} \boldsymbol{I}_O \Pr\{B, C\} \boldsymbol{I}_B \Pr\{G\} \sum_Q \Pr\{N, P, Q\} \boldsymbol{I}_Q.
\end{aligned} \tag{21}$$

With the initial condition $\Lambda(B,C,G,O,P) = \boldsymbol{I}_B \boldsymbol{I}_O$, the pivot border is now informed. By Corollary 9, $\Pr\{P,[B,O,Q]\}$ can be calculated as

$$\begin{aligned}
&\sum_{\{B,C,G,O\}} \Pi(B,C,G,O,P)\,\Lambda(B,C,G,O,P) \\
&= \sum_{\{B,C,G,O\}} \sum_{N} \Pr\{O|C,G,N\}\,\boldsymbol{I}_O \Pr\{B,C\}\,\boldsymbol{I}_B \Pr\{G\} \sum_{Q} \Pr\{N,P,Q\}\,\boldsymbol{I}_Q.
\end{aligned} \tag{22}$$

It is easy to verify that the RHS is indeed $\Pr\{P,[B,O,Q]\}$.

On the other hand, suppose we pick root $\{N,P,Q\}$ as the pivot border, then the collection phase includes:

1. By Theorem 27(b),
$$\Lambda(B,C,G,O,P) = \boldsymbol{I}_B \boldsymbol{I}_O.$$

2. By Lemma 23(b),
$$\begin{aligned}
\Lambda(B,C,G,N,P) &= \sum_{O} \Pr\{O|C,G,N\}\,\boldsymbol{I}_O \Lambda(B,C,G,O,P) \\
&= \sum_{O} \Pr\{O|C,G,N\}\,\boldsymbol{I}_O \boldsymbol{I}_B.
\end{aligned}$$

3. By Theorem 25 and Lemma 24, with the boundary conditions $\Pi_{\{B,C,G,N,P\}}(B,C) = \Pr\{B,C\}\,\boldsymbol{I}_B$ and $\Pi_{\{B,C,G,N,P\}}(G) = \Pr\{G\}$,
$$\begin{aligned}
&\Lambda(N,P) = \Lambda_{\{B,C,G,N,P\}}(N,P) \\
&= \sum_{\{B,C,G\}} \Pi_{\{B,C,G,N,P\}}(B,C)\,\Pi_{\{B,C,G,N,P\}}(G)\,\Lambda(B,C,G,N,P) \\
&= \sum_{\{B,C,G\}} \Pr\{B,C\}\,\boldsymbol{I}_B \Pr\{G\} \sum_{O} \Pr\{O|C,G,N\}\,\boldsymbol{I}_O.
\end{aligned}$$

4. By Theorem 25 and Lemma 23(b), at the pivot border,
$$\Lambda(N,P,Q) = \Lambda(N,P). \tag{23}$$

With the initial condition $\Pi(N,P,Q) = \Pr\{N,P,Q\}\boldsymbol{I}_Q$, the pivot border is now informed. By Corollary 9, $\Pr\{P,[B,O,Q]\}$ can be found by

$$\sum_{\{N,Q\}} \Pi(N,P,Q)\,\Lambda(N,P,Q)$$
$$= \sum_{\{N,Q\}} \Pr\{N,P,Q\}\boldsymbol{I}_Q \sum_{\{B,C,G\}} \Pr\{B,C\}\boldsymbol{I}_B \Pr\{G\} \sum_{O} \Pr\{O|C,G,N\}\boldsymbol{I}_O,$$

which is the same as Equation (22).

6.3.7 The Distribution Phase

In the collection phase, the messages converge to the pivot border; so it helps to know all the evidence variables and construct the border EC before the message propagations. Introducing a new evidence may require re-calculating some messages already sent to the pivot border. On the other hand, as discussed in §5.6, the distribution phase starts when the pivot border becomes the single member of the informed set $\mathcal{J}$. Introducing a new query border $\mathcal{Q}$, which includes a query variable, only requires calculating the messages from the corresponding gate in $\mathcal{J}$ to $\mathcal{Q}$, making $\mathcal{Q}$ informed and thus allow $\Pr\{\mathcal{Q},[\mathcal{E}\backslash\mathcal{Q}]\}\boldsymbol{I}_{\mathcal{Q}}$ to be calculated as $\Pi(\mathcal{Q})\Lambda(\mathcal{Q})$ by Theorem 8. There is no message re-calculation when the query variables are considered one-at-a-time.

For the BN $\mathbb{C}$ with $\mathcal{E} = \{B,O,Q\}$, assume we wish to obtain the posterior marginal of the outside variable M. With $\Lambda(N,P,Q)$ obtained in Equation (23), we can use Lemma 23 to send an upward message to border $\{N,M,Q\}$ as

$$\Lambda(N,M,Q) = \sum_{P} \Pr\{P|N,M,Q\}\boldsymbol{I}_Q\Lambda(N,P,Q)$$

Finally, with $\Pi(N,M,Q) = \Pr\{N,M,Q\}\boldsymbol{I}_Q$,

$$\Pr\{M,[B,O,Q]\} = \sum_{\{N,Q\}} \Pi(N,M,Q)\,\Lambda(N,M,Q) = \sum_{\{N,Q\}} \Pr\{N,M,Q\}\boldsymbol{I}_Q\Lambda(N,M,Q).$$

On the other hand, if we wish to obtain the posterior marginal of the outside variable F, then we can start with $\Pi(B,C,G,O,P)$ calculated in Equation (21):

1. Using Lemma 20,

$$\Pi(B,C,O,P) = \sum_{G} \Pi(B,C,G,O,P).$$

2. Using Lemma 20, $\Pi(B,F,O,P) = \sum_C \Pr\{F|B,C\}\, \boldsymbol{I}_B \Pi(B,C,O,P)$.

Finally, with $\Lambda(B,F,O,P) = \boldsymbol{I}_B \boldsymbol{I}_O$,

$$\Pr\{F,[B,O,Q]\} = \sum_{\{B,O,P\}} \Pi(B,F,O,P)\, \boldsymbol{I}_B \boldsymbol{I}_O.$$

If all non-evidence variables are the query variables, the messages may be sent "asynchronously," without a goal; that is, once a border becomes informed, it may send the messages to all of its neighbors that have not received a message from it. (See Díez & Mira, 1994.) Especially in this case,

1. If the number of children of $\mathcal{B}_i$ is large, then it may be advantageous to calculate $\Lambda(\mathcal{B}_i) = \prod_{\mathcal{W}\in\mathcal{L}_{\mathcal{B}_i}} \Lambda_{\mathcal{W}}(\mathcal{B}_i)$ first, then, for all $\mathcal{B}_{i+1,j} \in \mathcal{L}_{\mathcal{B}_i}$ in Theorem 22, we calculate $\Lambda(\mathcal{B}_i)/\Lambda_{\mathcal{B}_{i+1,j}}(\mathcal{B}_i)$ instead of $\prod_{\mathcal{W}\in\mathcal{L}_{\mathcal{B}_i}\setminus\mathcal{B}_{i+1,j}} \Lambda_{\mathcal{W}}(\mathcal{B}_i)$;

2. If the number of parents of $\mathcal{B}_i$ is large, then it may be advantageous to calculate $\Pi(\mathcal{H}_{\mathcal{B}_i}) = \prod_{\mathcal{W}\in\mathcal{H}_{\mathcal{B}_i}} \Pi_{\mathcal{B}_i}(\mathcal{W})$ first, then, for all $\mathcal{B}_{i-1,k} \in \mathcal{H}_{\mathcal{B}_i}$ in Lemma 24, we calculate $\Pi(\mathcal{H}_{\mathcal{B}_i})/\Pi_{\mathcal{B}_i}(\mathcal{B}_{i-1,k})$ instead of $\prod_{\mathcal{W}\in\mathcal{H}_{\mathcal{B}_i}\setminus\mathcal{B}_{i-1,k}} \Pi_{\mathcal{B}_i}(\mathcal{W})$.

7 Discussions

1. Our algorithm can handle what are known as the "soft evidences." An evidence variable E is soft if $\mathrm{Va}^e(E)$, its set of observed values, may have more than one members. In other words, the evidence indicator column $\boldsymbol{I}_E$ may have more than one non-zero values. This is less restrictive than the "hard evidence" assumption normally found in other inference algorithms, which requires $\mathrm{Va}^e(E)$ to have only one value. (See Langevin & Valtorta, 2008.)

2. All junction-tree based inferences share a worst-case complexity, which is exponential with respect to the largest clique size of the underlying undirected graph. According to Wu and Butz (2005), "Lauritzen and Spiegelhalter... were concerned with the size of the clique in the junction tree (transformed from the DAG of a BN), and they realized that their method would not be computational *feasible* if a large clique is present in the junction tree. The Hugin architecture has the same concern as the Lauritzen-Spiegelhalter

architecture, namely, the size of the clique in a junction tree. The Shafer-Shenoy architecture... used hypertree and Markov tree (junction tree) to describe the architecture. In [9], it was repeatedly emphasized that the efficiency and feasibility of their architecture depends on the size of the clique in a junction tree." ([9] referred to Shafer, 1996.) Wu and Butz (2005) then showed that "the presence of a node with a large number of parents can occur in both singly connected and multiply connected BNs. Therefore, in both singly and multiply connected BNs, the computation for exact inference will be exponential in the worst case."

In our algorithm, let us similarly assume that the largest border size in a border polytree is not too large. This imposes some limitations on the sizes of all cohort probability tables $\Phi(\cdot)$, which is dependent to the numbers of parents and children, and the number of possible values of each variable $V \in \mathcal{V}$.

3. The collection phase collects information about the evidences in the entire border polytree. It is intuitive that this process should be reduced to within the sub-polytree border EC: We only need to pass messages within it, toward its inside pivot border, starting from its evidential roots and leaves as in Theorem 27. In other words, in the collection phase, the BP is "pruned" to its border EC. As illustrated above with the BN $\mathbb{C}$ having $\mathcal{E} = \{B, O, Q\}$, the collection phase to the pivot border $\{N, P, Q\}$ only requires four message propagations within its border EC.

This message passing reduction is possible in our algorithm because we know the boundary conditions: Having a *directed* border polytree, we can use Theorem 26 because we can determine whether a border's outside neighbor is its parent or child. Also, it is important that all prior marginals $\Pr\{\mathcal{B}_i\}$ are pre-calculated (only once) and *pre-loaded* for this theorem. For example, to use Equation (21), besides the CPTs $\Pr\{O|C, G, N\}$ and $\Pr\{G\}$, we need the prior marginals $\Pr\{N, P, Q\}$ and $\Pr\{B, C\}$.

In many applications, the difference between the prior and the posterior marginals of a variable is more telling than the posterior itself, so pre-calculating the prior marginals should be done anyhow. Off-line, we calculate $\Pr\{\mathcal{B}_i\}$ as $\Pi(\mathcal{B}_i)$ without evidence, in a topological order of the macro-nodes and of the borders inside each macro-node, using Lemma 20 (with the complete cohort probability tables $\Phi(\cdot)$) and Lemma 21 (with $\Pi_{\mathcal{B}_i}(\mathcal{B}_{i-1,j}) = \Pi(\mathcal{B}_{i-1,j})$ by Theorem 22).

In the collection phase, each border in a EC sends only one message toward the pivot border. Thus the time complexity of the collection phase is lin-

ear with respect to the number of borders within the sub-polytree border EC. With one additional evidence, this number increases by the number of borders between it and the pivot border. If N is the sole evidence in the BN $\mathbb{C}$, then the collection phase is not needed, as the pivot border $\{N, P, Q\}$ is automatically informed with $\Lambda(N, P, Q) = \boldsymbol{I}_N$ and $\Pi(N, P, Q) = \Pr\{N, P, Q\}\boldsymbol{I}_N$, where $\Pr\{N, P, Q\}$ is pre-loaded. Thus the time complexity in the collection phase is linear with respect to the number of evidence variables.

4. Similarly, in the distribution phase, we pass messages from the growing informed set $\mathcal{J}$ to a query border $\mathcal{Q}$. As discussed above, the propagation starts from the informed gate connecting $\mathcal{J}$ with $\mathcal{Q}$; hence there is no need to re-visit any node inside $\mathcal{J}$. In other words, in the distribution phase, the BP is "pruned" to its "query core," which is the smallest sub-polytree that contains all the query borders. The computational complexity in the distribution phase is linear with respect to the number of borders inside the query core, which in turn is linear with the number of query variables. Again, the pre-loaded prior marginals $\Pr\{\mathcal{B}_i\}$ and the directed border polytrees are essential for the boundary conditions.

Consider the case where all non-evidence variables are the query variables. As far as we know, all junction-tree based inference architectures (including the LAZY propagation algorithm, Madsen & Jensen, 1999) require two passes through the entire network, one in each phase. Our algorithm requires one pass through the smaller border EC in the collection phase, and one pass through the entire network in the distribution phase. On the other hand, consider the case in which the single evidence variable and the single query variable are in the same border; our algorithm requires no message propagation in both phases.

5. In summary, the novel features in this paper are:

   i. The parentless polytree method (§6.1) to partition a BN into a macro-node polytree, by opening the loops in an otherwise growing parentless polytree.

   ii. The border algorithm to construct a directed chain from a BN (§2), or from a macro-node (§6.2).

Combining the above two algorithms, we can convert any Bayesian network into a border polytree. The border algorithm then provides the means to propagate the downward and upward messages in a border polytree, allowing us to calculate its posterior marginal probabilities (§6.3).

Also, the message propagations in the distribution phase is carried out one query border at a time, within the query core sub-polytree only (§5.5).

With the above novel features, *the time complexity of our inferences in a Bayesian network is linear with respect to the number of its evidence and query variables, regardless of the number of borders in its corresponding border polytree, or the number of its variables.*

References

1. Chang, K., and Fung, R. (1989) "Node Aggregation for Distributed Inference in Bayesian Networks" In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, Michigan, 265–270.

2. Cooper, G. F. (1990) "The computational complexity of probabilistic inference using Bayesian belief networks" *Artificial Intelligence,* 42, 393–405.

3. Darwiche, A. (2003) "A differential approach to inference in Bayesian networks" *Journal of the ACM*, 50, 280-305.

4. Dechter, R. (1999) "Bucket elimination: A unifying framework for reasoning" *Artificial Intelligence*, 113, 41–85.

5. Díez, F. (1996) "Local conditioning in Bayesian networks" *Artificial Intelligence*, 87(1996).

6. Díez, F. and Mira, J. (1994) "Distributed inference in Bayesian networks" *Cybernetics and Systems*, 25, 39–61.

7. Guo, H. and Hsu, W. (2002) "A survey of algorithms for real-time Bayesian network inference" In the joint AAAI-02/KDD-02/UAI-02 workshop on Real-Time Decision Support and Diagnosis Systems.

8. Kim, J. H. and Pearl J. (1983) "A computational model for combined causal and diagnostic reasoning in inference engines" *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, Karlsruhe, West Germany, 190–193.

9. Koller, D. and Friedman, N. (2009) *Probabilistic Graphical Models: Principles and Techniques,* Massachusetts: MIT Press.

10. Langevin, S. and Valtorta, M. (2008) "Performance evaluation of algorithms for soft evidential update in Bayesian networks: First results" *Scalable Uncertainty Management*, 284–297.

11. Lauritzen, S. L. and Spiegelhalter, D. J. (1988) "Local computations with probabilities on graphical structures and their applications to expert systems" *Journal of the Royal Statistical Society*, Series B, 50, 157-224.

12. Lepar, V. and Shenoy, P. (1999). "A comparison of Lauritzen-Spiegelhalter, Hugin, and Shenoy-Shafer architectures for computing marginals of probability distributions" in G. F. Cooper & S. Moral (eds.), *Uncertainty in Artificial Intelligence*, 14, 328-337, Morgan Kaufmann, San Francisco, CA.

13. Madsen, A. L. and Jensen, F. V. (1999) "LAZY propagation: A junction tree inference algorithm based on lazy evaluation" *Artificial Intelligence*, 113, 203-245.

14. Ng, K. and Levitt, T. S. (1994) "Incremental Dynamic Construction of Layered Polytree Networks" *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence, 440-446.*

15. Pearl, J. (1986a) "A constraint-propagation approach to probabilistic reasoning" In L. N. Kanal and J. F. Lemmer (Eds), *Proceedings of the 2nd Conference on Uncertainty in Artificial Intelligence*, Amsterdam, NorthHolland, 357-369.

16. Pearl, J. (1986b) "Fusion, propagation and structuring in belief networks" *Artificial Intelligence*, 29, 241-288.

17. Pearl, J. (1987) "Evidential reasoning using stochastic simulation of causal models" *Artificial Intelligence*, 32, 245–257.

18. Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann.

19. Peot, M. A., and Shachter, R. D. (1991) "Fusion and propagation with multiple observations in belief networks." *Artificial Intelligence*, 48, 299–318.

20. Shachter, R. D. (1990) "Evidence absorption and propagation through evidence reversals" In M. Henrion, R. D. Shachter, J. F. Lemmer, and L. N. Kanal (Eds.), *Uncertainty in Artificial Intelligence,* 5, 173-190.

21. Shachter, R. D., D'Ambrosio, B., Del Favero, B. D. (1990) "Symbolic probabilistic inference in belief networks" *Proceedings of the 8th National Conference on Artificial Intelligence*, MIT Press, Boston, 126-131.

22. Shafer, G. (1996) *Probabilistic Expert Systems*, Society for Industrial and Applied Mathematics.

23. Wu, D. and Butz C. "On the complexity of probabilistic inference in singly connected Bayesian networks" *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing*, 581-590.

24. Zhang, N. L. and Poole, D. (1996) "Exploiting causal independence in Bayesian network inference" *Journal of Artificial Intelligence Research*, 5, 301–328.